IBM Business Analytics and Cloud Computing

Best Practices for Deploying Cognos Business Intelligence to the IBM Cloud

Anant Jhingran
Stephan Jou
William Lee
Thanh Pham
Biraj Saha

IBM BUSINESS ANALYTICS AND CLOUD COMPUTING

Best Practices for Deploying Cognos Business Intelligence to the IBM Cloud

Anant Jhingran
Stephan Jou
William Lee
Thanh Pham
Biraj Saha

MC Press Online, LLC
Ketchum, ID 83340

IBM Business Analytics and Cloud Computing
Anant Jhingran, Stephan Jou, William Lee, Thanh Pham, and Biraj Saha

First Edition
First Printing—September 2010

MC Press Online, LLC
Corporate Offices
P.O. Box 4886
Ketchum, ID 83340-4886 USA
For information regarding sales and/or customer service, please contact:
MC Press
P.O. Box 4300
Big Sandy, TX 75755-4300 USA
For information regarding permissions or special orders, please contact:
mcbooks@mcpressonline.com

ISBN: 978-158347-363-4

Acknowledgments

Thanks to Ka Weng Mah and Gilbert Chan of Nanyang Polytechnic at Singapore for their assistance with the Cognos 10 installation and to Scott Crowder at IBM for his contributions to the section on cloud computing workloads.

We would also like to thank Christopher Dziekan, Dino Georgopoulos, Kevin Haaland, Tracy Harris, Keith Lindsey, Chris McPherson, Susumu Sai, Clint Sherwood, Fenny Thomas, Pierre Touissant, Ralf Vierich, the Hursley Java Information Manager team, and many other colleagues from IBM for their support, assistance, and reviews of the content in this book.

About the Authors

Dr. Anant Jhingran is an IBM Fellow and Vice President and Chief Technical Officer (CTO) for IBM's Information Management division. He is deeply involved in IBM's middleware in the cloud (including co-chairing IBM's Cloud Computing Architecture Board) and with the technical strategy of products and solutions for databases, content management, business intelligence, and information integration.

Before assuming his present role, Anant led the IBM team responsible for designing and building solutions to meet the requirements of business analytics on structured and unstructured data. He has also served as Director of Computer Science at the IBM Almaden Research Center and as Senior Manager for E-Commerce and Data Management at IBM's T. J. Watson Research Center. Anant received his Ph.D. from the University of California at Berkeley in 1990. He holds several patents and awards for DB2 and is a member of the IBM Academy of Technology.

Stephan Jou is a Technical Architect, Research Staff Member, and Senior Technical Staff Member in IBM's Business Analytics division in the Technology and Innovation group at the Office of the CTO. In his career at Cognos Software, he architected and led the development and commercialization of several initial-release products in the areas of data mining, neural networks, visualization,

mobile, dashboarding, and semantic search. His current role at IBM focuses on translating academic and IBM research into product strategy for Cognos and SPSS Software.

Stephan holds a Master of Science in Computational Neuroscience and Biomedical Engineering as well as dual Bachelor of Science degrees in Computer Science and Human Physiology, all from the University of Toronto.

William Lee is a Senior Software Consulting Engineer at IBM through the Cognos acquisition. He is a member of the Technology and Innovation team for the Office of the CTO in IBM's Business Analytics division, where he helps to define the technical vision and direction for Cognos and SPSS Software products.

William has been with Cognos and IBM since 1992. He holds a Bachelor of Computer Science and Mathematics and a Master of Computer Science from Carleton University, Ottawa, Canada.

Thanh Pham is a Solution Architect in the IBM Information Management Advanced Technology group. His current focus is on helping customers build applications using the IBM Mashup Center product and IBM cloud computing. Before this role, he was an Architect for the ECM/Filenet® Business Process Framework.

Biraj Saha is an Advisory Software Developer at IBM Cognos, specializing in metadata and algorithm design and development for Cognos modeling tools such as Framework Manager, Metrics Designer, and Architect, as well as SOA and SDK development for Cognos BI Server. He was formerly a Senior Software Engineer for EDS Systemhouse, serving in lead development roles for a wide array of customers on various RDBMS-related developments, including ERP and RDBMS-vendor application conversions as well as custom Java, C++, stored procedure, and 4GL applications.

Biraj holds a Bachelor of Computer Science from the University of New Brunswick, Canada, and a Master of Computer Science, specializing in object-oriented database constraint theory, from the University of Waterloo, Canada.

Table of Contents

Introduction

This book is intended to be a practitioner's guide to IBM® Cognos™ software and the cloud. If you are a solution architect or IT architect and have been charged with evaluating, designing, or implementing a Cognos Business Intelligence solution to leverage a cloud infrastructure, this book is for you.

Rather than spend much time on abstract notions of cloud computing, this book focuses on pragmatic and practical information on the best practices, guidelines, and even specific software and configuration steps to help you achieve success with your Cognos analytics cloud solution.

The Organization of This Book

This book is made up of seven chapters:

- Chapter 1 gives a brief introduction to cloud computing, workloads in cloud computing, analytics as a service, and its implementation using Cognos software.
- In Chapter 2, you start designing your cloud solution at a high level. This chapter provides guidance and considerations for your cloud solution, such

as where the data should be located and how many images you should create.

- Chapter 3, the longest chapter, provides detailed, step-by-step instructions on how to install, configure, secure, and test a Cognos Business Intelligence deployment in the IBM Cloud. The result of this chapter is a saved private instance to an image that you can then deploy in minutes, on demand. While the scenario used in this chapter results in a single image, the techniques described form the basis to produce more complicated, multiple-image topologies.
- Chapters 4 to 7 address various considerations once you graduate from single images to multiple images and more complicated topologies that are better suited for production workloads:
  - » Chapter 4 focuses on the important topic of securing your Cognos cloud solution and provides some best practices.
  - » Chapter 5 describes techniques that can be used when moving from a single image to multiple images. You can use these techniques to dynamically connect and disconnect instances in your Cognos Business Intelligence solution.
  - » Chapter 6 discusses scalability and performance, how to plan for scaling, and how to ensure a high-performing solution.
  - » Chapter 7 focuses on high availability, providing tips to ensure that your Cognos BI cloud solution is highly available and handles failover scenarios.

Conventions Used in This Book

The following conventions are used in this book:

- New terms are indicated in *italic*.
- Commands that should be entered from a Unix® prompt are indicated with a monospaced font and prefixed with a dollar sign:

```
$ Unix Command
```

- Commands that should be entered from a Microsoft® Windows® command prompt are indicated with a monospaced font and prefixed with a greater-than symbol (>):

```
> Windows Command
```

- Output from the computer is indicated with a monospaced font:

```
Text Output
```

- Commands containing words that should be customized for your environment appear in **bold**:

```
$ Unix Command with words that should be customized
```

- If a command occurs on a single line but does not fit across the width of a page in this book, the text will wrap to an indented second line:

```
$ This is a Unix command that should be entered on a single line,
  but it is so long that we need to print it on two lines, with
  the subsequent lines indented.
```

- File and directory names, system names, program names, and similar elements that are embedded in the regular text of a paragraph appear in a **special font**.

Corrections and Errata

Any errors or corrections discovered after publication of this book will be posted on the Web at *http://www.jou.ca/cognoscloudbook* and on the book's page at *http://www.mc-store.com*.

1

Cloud Computing and Analytics

At the time of writing, a Google search on the phrase "cloud computing definition" returned more than 3.5 million results. There appear to be as many definitions of cloud computing as there are people excited about it! Some of these definitions are very good. For example, the U.S. National Institute of Standards and Technology (NIST) provides a concise but comprehensive effort at *http://csrc.nist.gov/groups/SNS/cloud-computing*.

This book will not repeat such efforts at defining cloud computing. Instead, we intend this book to be a practical companion to leveraging cloud computing in your IBM Cognos analytics solution. As such, it focuses on the main characteristics of cloud computing with respect to their tangible advantages for you, the cloud practitioner.

On-Demand Infrastructure

On cloud computing platforms, the required IT infrastructure for your applications is provided to you, based on what you actually require. Nearly all clouds now provide compute cycles, networking, storage space, and memory capacity, all on an on-demand basis. Because you can simply release unused resources back into the pool, you do not have to worry about over-purchasing more hardware than you actually need.

In a pure and simple comparison with traditional data centers, this arrangement provides immediate and obvious cost advantages. Underutilization of purchased hardware is a genuine problem. It's what made virtualization such an attractive IT strategy in the early 2000s: replace the physical hardware with virtual hardware so you can allocate virtual machines when you need them and deallocate them when you're done. This strategy is particularly cost-effective for analytical applications that are tied to seasonal behavior, such as a sales application that is used only during the end of a quarter.

Small wonder that the major cloud platforms, including those from Amazon and IBM, are, at their lowest level, Web interfaces wrapped around virtual machines (VMs), storage, and networking. Being able to create and configure VMs through a simple browser interface or through Representational State Transfer (REST) calls is one simple way to think about and approach cloud computing.

This pay-as-you-go, utility-based cost model is, in some ways, the most innovative aspect of cloud computing. You trade away the requirement for up-front capital expenditures (capex) to purchase hardware and software, and instead favor ongoing operational expenditures (opex) based on what you actually use.

This book takes you through the process of leveraging such an infrastructure to create a fully working IBM Cognos Business Intelligence (BI) virtual instance running in the cloud. This instance is then saved as an image, consuming no resources or cost, until you are ready and have a need for a Cognos deployment.

While the steps in this book are based on the IBM Smart Business Development and Test Cloud, they are also applicable with little modification to other cloud infrastructures, such as Amazon's. And of course, the best practices we describe here have general applicability and relevance, no matter how you ultimately deploy your Cognos application.

On-Demand Higher Services

Moving above the so-called Infrastructure-as-a-Service (IaaS) layer to the higher so-called Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS) layers is where cloud starts to differentiate itself from simple virtualization. The PaaS and SaaS cloud layers bring higher-level services to the table, and things get

much more interesting. Rather than thinking in terms of machines, networking, bandwidth, and storage space, imagine services related to the provisioning of complex topologies with defined quality-of-service constraints, analytical and reporting services, and Hadoop-style big-calculation jobs.

While nascent, there is a tremendous amount of growth in the cloud computing ecosystem around these higher-level cloud services. For example, both Amazon and Yahoo! offer platforms that can execute Hadoop applications in their cloud infrastructures.

There is a good and practical reason why higher-level services are interesting: they are more cost-efficient. For example, running a Hadoop job using Amazon's Elastic MapReduce, which leverages Amazon's Elastic Compute (EC2) and Simple Storage Service (S3) under the covers, costs less than directly using EC2 or S3 yourself. That's because not only do you avoid having to install and configure the software, but Amazon can optimize and manage the entire infrastructural stack much more efficiently.

The end result is that as we move up the cloud stack and focus more on higher-level services that provide targeted solutions and workloads, we are able to build more for less.

Resource Pooling and Rapid Elasticity

The distressing amount of hardware underutilization in traditional data centers that we noted previously remains the main reason why virtualization and cloud has been on the IT agenda for the past few years and will continue to be in the years to come. Being able to more closely match capacity and cost with demand is the cost justification we all need.

This ability to pool and share resources to match demand clearly requires rapid elasticity. The amount of storage available (and being paid for) should always be slightly ahead of the growth of your database. Any new virtual machines required to handle added load should be recruited and connected to the system in minutes or hours, not days or weeks.

This book provides techniques and best practices to scale up or down a Cognos BI system, dynamically recruiting additional virtual machines and connecting them or disconnecting them as required.

Flexible Deployment Models

When most people first hear of cloud computing, they think of the public cloud—an IT infrastructure that is delivered externally. While a public cloud is appropriate for many scenarios, there are many other cases in which data cannot leave the enterprise boundaries. Fortunately, in addition to public clouds, you can deploy your solution to private clouds, where the cloud infrastructure is erected within the enterprise firewalls and managed by the enterprise IT department itself, or even to hybrid clouds, which are a combination of public and private clouds, with systems on both sides and a secure connection between them.

No single model will work in all cases. Data security and sensitivity, bandwidth and latency, and even legal and regulatory requirements all need to factor into the deployment and topology of this solution. Fortunately, cloud computing offers a flexible allocation of resources and the ability to loosely couple systems together with standard Internet protocols, letting you design the right solution for the requirements on hand.

This book provides some general guidance about which components of your Cognos deployment can be located in the cloud, and when.

The Workload Model for Cloud

When you do any sort of reading in the area of cloud computing, you quickly run into the concept of a *workload*. A workload is a set of operations executed on IT resources for a particular purpose and typically considered as a single logical element. The key components of a given workload are the application(s), the usage pattern, the service level agreement, and a data structure. For example, you might hear people talk about a "departmental BI with a seasonal transaction model but large data volume" workload.

A workload approach to cloud computing helps us understand when an application is ideal for the cloud, as well as which cloud architecture is most appropriate. IBM has analyzed the various workload types, based on the relative cost and benefits of leveraging cloud computing. Figures 1.1 and 1.2 summarize the results for a typical example in each broad workload category for external and private clouds.

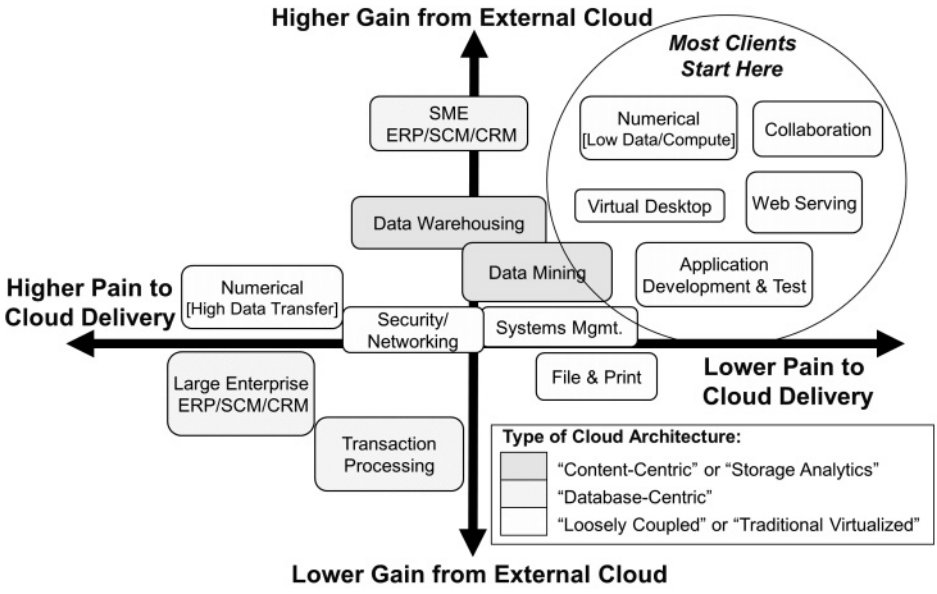

Figure 1.1: The relative affinity of different workload categories for public or external clouds.

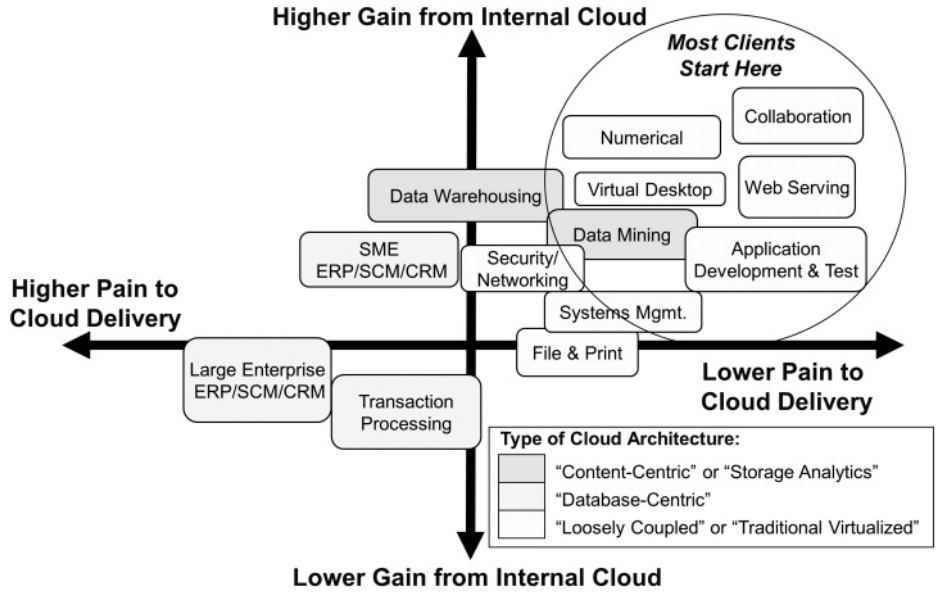

Figure 1.2: The relative affinity of different workload categories for private or internal clouds.

These figures serve as good guidance for representative workloads within each category, but, of course, there are always unique situations in each category. A more detailed cloud affinity tool that lets you input the specific characteristics of your particular analytics workload (or other workload types) is available at *http://freedom.researchlabs.ibm.com/ibmappcr/applications/p1/CloudAffinityAnalyzer/CloudAffinityAnalyzer.html*.

An analytical workload needs to be examined first to see how appropriate it is for the cloud and then for what type of cloud deployment makes sense. This book covers several considerations, but the main ones are usually the nature of the data being analyzed and its sensitivity to public exposure. If the data is too large or too sensitive to be moved, a private or possibly a hybrid cloud is a strong candidate.

Examples of ideal IBM Cognos Business Intelligence workloads for the IBM Smart Business Development and Test Cloud include the following:

- Development and test workloads, including pilots and proofs-of-concept, which typically involve non-sensitive or small amounts of data
- Standalone BI implementations
- Variable or seasonable workloads that take advantage of peaks and valleys and load balance between on-premises and cloud-based systems
- Cloud-to-cloud applications, where the data is coming from another cloud service

These workloads all take advantage of several key benefits that cloud computing provides:

- Significant hardware and software cost savings
- The flexibility of opex versus capex
- On-demand, elastic IT resources
- Faster time to provisioning

These benefits enable you to standardize and share costs while maintaining the control and ownership within your IT department.

This book, and particularly Chapter 2, provides best practices on other characteristics of your analytics workload and its implications for your Cognos Business Intelligence solution architecture.

Cognos Software and Analytics as a Service

The concept of analytics as a service is a compelling one: identify higher-level services associated with analytical functions—such as reporting, querying, prediction, and exploration—and provide those in a cloud-hosted model. As we discussed earlier, providing such higher-level services in the cloud can have several advantages. Although it's still early days in this area, there are already a number of interesting and innovative analytics-as-service experiments in the cloud ecosystem.

A cloud deployment of IBM Cognos BI actually provides much of what you might expect from an analytics as a service offering and gives a hint of what might be in the future. With a platform built on service-oriented architecture concepts, Cognos 8 and Cognos 10 BI provide programmatic interfaces that let their capabilities be leveraged by both Cognos and custom applications. There are several integration points:

- Cognos portlets can be embedded in external applications through the portlet SDK, Web Services for Remote Portlets (WSRP), Webparts, or iWidgets.
- Web content can be embedded in Cognos portlets or reports through the portlet SDK or HTML components.
- Reports and data can be processed by external applications, using the SOAP SDK or Cognos Mashup Service.
- Data, models, and other content can be inserted or manipulated through the SOAP SDK, FM API, or data drivers.

Many of the relevant application program interfaces in these integration points are cloud-friendly, leveraging standard protocols such as ATOM, REST, SOAP, and XML. As a result, integrating to your cloud-deployed instance of IBM Cognos Business Intelligence is often an exercise in ensuring that you have network connectivity between the client machine and your cloud instance and then pointing the client application to the instance's dispatcher URI.

This book provides steps on how to configure other IBM Cognos applications that do not run in the cloud (e.g., Framework Manager, Mobile, Office, PowerPlay, and Transformer) and point them to your cloud installation of the Cognos Business Intelligence Server. This background, along with the details provided by the various Cognos SDK manuals, will ensure that you, too, will be successful in using your own Cognos analytics cloud deployment as a service.

2

Getting Started

The first steps in getting started with cloud computing involve data, security, topology, and Linux considerations.

Data Considerations

Databases and data sources can be co-located in the cloud with your cloud application. Or, they can be located on-premises behind your firewall, but with a secure connection to your cloud. Cognos Business Intelligence has two classes of data sources: 1) the content store and metric store database and 2) the query databases and other data sources. Figure 2.1 depicts these aspects of the Cognos 8 tiered architecture.

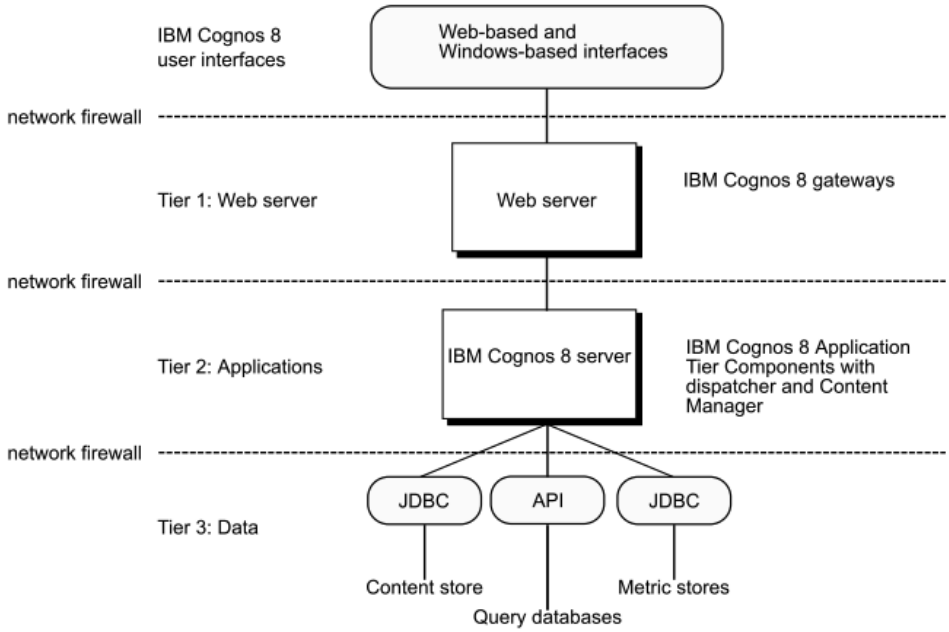

Figure 2.1: The Cognos 8 tiered architecture and data sources.

The *content store* is a relational database that contains data that Cognos BI needs, including report specifications, published models (and the packages that contain them), connection information for data sources, information about your users, and information about scheduling and bursting reports. The *metric store* is the equivalent of the content store for Metric Studio (an optional component of Cognos BI). It contains content for metric packages and other Metric Studio settings, such as user preferences. If you are not using Metric Studio, you do not need a metric store.

The *query databases* are relational databases that can be accessed through Cognos BI. They provide the data for its reports and analyses, through a JDBC or Virtual View Manager connection.

The *data sources* include all relational databases. Other, less common data sources can be accessed through Cognos BI, as well. These are not relational databases; they are things such as dimensional cubes and files.

For best performance, the content store and metric store databases should be as close as possible to the application on the network. Close proximity on the network minimizes latency between the Cognos BI Server components and the content and metric store databases. Ideally, therefore, these databases should be

in the IBM Cloud environment—either in a separate virtual machine instance or with your application tier components in the same instance.

For your query databases, consider the specifics of your intended workload and scenario, as outlined in Table 2.1. Perhaps, for example, your reports require high performance or rapid querying of your data, and this data can easily be moved to the cloud, too. At the same time, privacy or security concerns are not a priority. In this case, you can realize significant cost savings by creating the query databases in the IBM Cloud.

In other situations, high performance or rapid data querying is not a priority, and there is a large amount of data that is difficult or expensive to move. Privacy, security, or other legal reasons may require you to maintain the data within your corporate firewalls. In this case, the query databases should be kept on-premises, within the network bounded by your firewalls. These databases can then be accessed from the IBM Cloud through a secure network connection. This configuration is sometimes referred to as a "hybrid cloud" environment because it is a mix of cloud instances and traditional behind-the-firewall instances.

In some circumstances, your query database is already in the cloud (e.g., Salesforce data). In this case, the security and latency challenges associated with the query data are not new to a cloud solution. Such "cloud-born" query data sources are ideal candidates for leaving in the cloud.

Table 2.1: Considerations for Locating Cognos Query Databases

Query Database in the Cloud	Query Database On-Premises
Workloads require high performance (e.g., rapid queries or large amounts of data).	Workloads have acceptable query performance over a network connection.
New or existing data is easily moved to the cloud (e.g., test data) or is "cloud-born" (e.g., Salesforce data).	A large amount of data exists, or the data is difficult to move to the cloud.
An acceptable level of privacy/security comfort exists around the location of data (e.g., public or non-sensitive data).	Privacy, security, or legal reasons require data to remain on-premises.

Another combination is also worth mentioning. In some situations, database replication can be used to copy an on-premises database instance to a database instance in the cloud. For example, one of IBM DB2®'s various replication alternatives might be an attractive option for you, provided you leverage a secured connection for the transaction.

File-based data sources, such as dimensional cubes and other files, are usually amenable to synchronization or transport to the cloud in a directory that is

accessible to the Cognos application instance. They can also be synchronized to an IBM Cloud storage instance that appears as a mounted directory to your instance.

Security Provider Considerations

Many Cognos applications authenticate users through third-party security tools such as LDAP or Active Directory. Such third-party authentication sources are typically used to create groups of users and to restrict content access to certain user groups. If your workload includes such a requirement, your cloud topology will also need to include your authentication source, as Figure 2.2 illustrates.

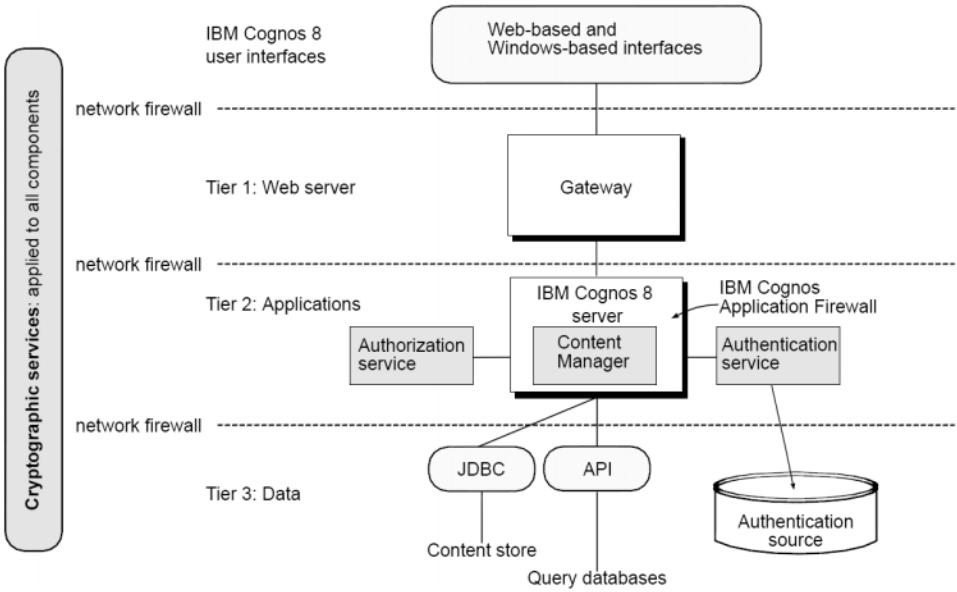

Figure 2.2: The Cognos tiered architecture with authentication.

In addition to data considerations, you will need to consider the location of your authentication source in this scenario (Table 2.2). You can co-locate your authentication source in the cloud, as you did with your Cognos application, or you can keep it on-premises behind your firewall with a secure connection to the cloud.

Table 2.2: Determining the Placement of the Authentication Source	
Security Provider in the Cloud	**Security Provider On-Premises**
Workloads require the best network latency or highest performance.	Workloads have acceptable query performance over a network connection.
User data is new or easily moved to the cloud (e.g., test-user data).	The user directory is too large or difficult to move to the cloud (as in the case of password synchronization).
An acceptable level of privacy/security comfort exists around the location of user data.	Privacy, security, or legal reasons require you to maintain user information on-premises.

Perhaps you're providing IT resources to a new or small set of users, such as a development or test/QC group. If the user data can be easily moved to the cloud (because it is small or static enough to be managed via SFTP or database replication) and there are no security or privacy concerns, then installing the authentication source in the cloud alongside or within your application cloud instance is recommended. On the other hand, a secured network connection between the on-premises authentication tool and the Cognos application can be used if any of the following are true:

- The authentication source needs to leverage an existing on-premises installation.
- The data cannot be easily replicated or moved.
- Privacy, security, or performance concerns exist.

Designing and Testing Your Topology

As you design and refine your topology, start simply and avoid unnecessary complexity. Satisfy your requirements, but always keep the number of cloud instances in your topology as low as possible. Adding instances is easy in the event that you encounter a future need based on additional requirements or load, so it's usually preferable to underestimate your requirements initially.

Starting simply also applies when it comes to the number of unique cloud images. For example, it is easier to manage a single DB2 database image that customizes itself on startup than to create five different query database images. Or, rather than have a Cognos content store image and a Cognos report server image, it might be simpler to have a generic Cognos BI image, with which you start only the services required for the instance upon startup.

The process of designing and then testing your topology is very iterative, likely consisting of these steps:

1. Design/refine your topology.
2. Create/customize the required instances.
3. Install/configure the instances.
4. Save image snapshots.
5. Test for functionality and performance.
6. Repeat.

If you start simply and keep your topology as small as possible, each of these steps becomes as straightforward as possible. We offer other tips in the pages that follow to help reduce the number of unique images required for your solution.

Embracing Linux®

The IBM Cloud gives you the flexibility to choose from several different operating systems on your instances, including Microsoft Windows and several versions of Linux. There are advantages to opting for one of the Linux-based distributions:

- The underlying infrastructure of the IBM Cloud is based on Linux technologies, including virtualization based on Xen.
- The licensing costs associated with Linux are lower than for Windows-based operating systems.
- Many cloud-appropriate applications and technologies—and the associated user communities and knowledge bases—are already focused on Linux environments.

Other things being equal, embracing Linux can thus be a very efficient and cost-effective part of your cloud strategy. Within your cloud design, deployment, and operations teams, there is also probably a rich storehouse of Linux experience and skills, which you can leverage to ensure the success of your project.

3

Installation and Configuration

This chapter describes the steps required to install and configure IBM Cognos Business Intelligence, either Version 8 or Version 10, into a single IBM Smart Business Development and Test Cloud image. When finished, you will have created a private image on the IBM Cloud. You can use this image to create and deploy instances of a fully configured and operational Cognos 8 or Cognos 10 instance, which can immediately be used for your analytics applications.

The techniques and information used to create this single image of Cognos 8 or Cognos 10 will form the basis for extending out to multiple image topologies. In addition, while the steps assume IBM DB2 as the database, IBM WebSphere® Application Server as the application server, and IBM HTTP Server as the Web server, the information should be useful to whatever configuration you are ultimately aiming for.

Set Up the Windows Client

While the server instances will be in the IBM Cloud, several tools allow easy access and management of the cloud instances from a PC running Microsoft Windows. Installing these tools is a one-time step; after that, you can use them to connect to any instance created in the IBM Cloud.

Install SSH Client (PuTTY)

You need a Secure Shell (SSH) client to log into your newly created instance and interact with it. If you are using a Windows machine, we recommend the PuTTY freeware SSH client. (The latest PuTTY version as of June 2010 is 0.60.)

Install WinSCP

You also need an SSH file transfer program to move software and data to your cloud instances. For Windows machines, we recommend WinSCP freeware. (The latest WinSCP version as of August 2010 is 4.2.8.)

Install X-Windows

To run X-Windows programs from the cloud, such as the Cognos Configuration tool, and have them appear on your machine, you will need an X-Windows client. On Windows machines, you can choose to install a commercial product, such as OpenText's Exceed, or opt to use one of several freeware clients, such as Xming.

If you are using Xming, follow these steps:

1. Download and install Xming, following the instructions available from the download site. (The latest Xming version as of August 2010 is 7.5.0.24.)
2. Launch Xming, and confirm that the Xming icon appears in the system tray.

Set Up and Configure the Cloud Instance

This section summarizes the steps involved to create your machine instance in the cloud. More detailed documentation is available online on the IBM Smart Business Development and Test Cloud documentation areas and asset catalog.

Create the Cloud Instance

To create the cloud instance, take these steps:

1. Browse to *http://www.ibm.com/cloud/enterprise*.
2. Register for an account, if required.
3. Sign in with your user ID and password.
4. In the Control Panel, click the **Add Instance** button.

5. In the "Add Instance" image selection that follows, scroll down and select "Red Hat Enterprise Linux 5.4 (32-bit)" to choose a Red Hat Enterprise Linux 5.4 base operating system instance. (For a 64-bit install, select "Red Hat Enterprise Linux 5.4 64-bit.") Click **Next**.
6. Specify your desired instance name and size type (Medium or Large), and then click **Add Key**.
7. Specify the name of the key to generate, and click **Continue**.
8. Ensure you download the private key file and save it. Do this by clicking **Save** or by creating the key manually. To create the key manually:
   a. Choose to open the file during download. Your key will appear in a separate Web page, with a **BEGIN RSA PRIVATE KEY** line at the top of the file and an **END RSA PRIVATE KEY** line at the bottom.
   b. Copy and paste the page contents (including the **BEGIN** and **END** lines) into a new file using a text editor.
9. Save the file with a name such as **ibmcloud_your.name@ibm.com_rsa**. Select your new key.
10. Verify overall selections for the new instance, and click **Next**.
11. Read and agree to the terms and conditions. Click **Submit**.
12. The Add Instance dialog's "Submitting" dialog appears. When the Add Instance dialog reports "Success," click **Close**.
13. The Control Panel should now list the new instance with a status of "Active," as shown in Figure 3.1.

Figure 3.1: The Control Panel showing an active instance.

Modify Security Permissions

In a single-image environment, you want to ensure that the permissions are correctly set up so that the **idcuser** account is able to install and run all the

software. We describe this process below. Keep in mind that there will be accounts with least privileges accessing the instance. Because this is the case, consider revoking root privileges after installation, or create a separate user account just to run Cognos 8 or Cognos 10.

1. Change the group for **idcuser** to **root** by setting the value of the fourth field to 0 (zero):

```
$ sudo vi /etc/passwd
idcuser:x:500:0:idcuser:/home/idcuser:/bin/bash
```

2. Prepare the target root installation folder and set group-level write permissions:

```
$ sudo mkdir -p /opt/ibm
$ sudo chmod -Rf g+w ibm
```

This step is important for the Cognos installation program to succeed.

Enable X11 Forwarding on Your Cloud Instance

X11 Forwarding is the part of X-Windows that permits messages from the cloud instance's Windows manager to be "forwarded" to your Windows client. By default, your instance should be set up for X11 Forwarding, but we will verify that here.

1. As a root user, ensure that the ports required for X11 forwarding are enabled:

```
$ sudo /sbin/iptables -L
Chain INPUT (policy DROP)
target    prot opt source       destination
ACCEPT    all  --  anywhere     anywhere     state RELATED,ESTABLISHED
ACCEPT    tcp  --  anywhere     anywhere     tcp dpt:ssh
Chain FORWARD (policy DROP)
target    prot opt source       destination
Chain OUTPUT (policy ACCEPT)
target    prot opt source       destination
```

2. Confirm that all policy values are "ACCEPT." If this is not the case (as shown above), you will need to update the values accordingly in the main **iptables** configuration file, **/etc/sysconf/iptables**, and then restart the **iptables** service:

```
$ sudo vi /etc/sysconfig/iptables
```

The file should look something like this:

```
# Generated by iptables-save v1.3.5 on Thu Mar  4 22:07:58 2010
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p tcp -m tcp --dport 22 -j ACCEPT
COMMIT
# Completed on Thu Mar  4 22:07:58 2010
```

3. Restart the service:

```
$ sudo /sbin/service iptables stop
$ sudo /sbin/service iptables start
```

Install X11 Support in Your Cloud Instance

In addition to the X-Windows client, you'll need to install Motif on the cloud instance to run X-Windows programs such as the Cognos Configuration tool on the cloud.

1. Create a temporary directory, such as **/home/idcuser/installs/centos**.
2. Add the CentOS repository to your system by creating the file **Centos.repo** in that directory, or create it on your local machine and use WinSCP to transfer the file to that directory (this step might not be required if your RedHat image is activated and connected to the RedHat repositories; in that case, skip to step 4):

```
[base]
name=CentOS-$releasever - Base
mirrorlist=http://mirrorlist.centos.org/?release=5.5&arch=x86_64&re
  po=os
#baseurl=http://mirror.centos.org/centos/$releasever/os/$basearch/
gpgcheck=1
gpgkey=http://mirror.centos.org/centos/RPM-GPG-KEY-CentOS-5
protect=1
[updates]
name=CentOS-$releasever - Updates
mirrorlist=http://mirrorlist.centos.org/?release=5.5&arch=x86_64&re
  po=updates
#baseurl=http://mirror.centos.org/centos/$releasever/
  updates/$basearch/
gpgcheck=1
gpgkey=http://mirror.centos.org/centos/RPM-GPG-KEY-CentOS-5
protect=1
```

3. Update your system to recognize the CentOS repository:

```
$ sudo cp /home/idcuser/installs/centos/Centos.repo /etc/yum.
  repos.d/
```

4. Download OpenMotif and associated X-Windows libraries:

```
$ sudo yum install openmotif.i386 xorg-x11-xauth libXtst xterm
$ sudo ln -s /usr/lib/libXm.so.4 /usr/lib/libXm.so.3
```

5. Edit **/etc/ssh/sshd_config** and ensure that the following lines are uncommented:

```
X11Forwarding yes
X11UseLocalHost yes
AllowTcpForwarding yes
```

6. If changes to **sshd_config** were required, restart the **sshd** service:

```
$ sudo /sbin/service sshd restart
```

Configure the Windows Client

Now that you have created the cloud instance, follow the steps described below to connect the Windows client to it. You can repeat these steps for all the IBM Cloud instances that are available to you.

Configure PuTTY

For each cloud instance, create a PuTTY session that lets you connect to it:

1. PuTTY requires a "ppk" version of the private key previously downloaded from the IBM Cloud. This step is required only once for each keyfile you have from the IBM Cloud.
   a. Launch **puttygen.exe**.
   b. Click **Load**, and choose the keyfile (e.g., **ibmcloud_your.name@ibm.com_rsa**). Click **OK** when you see the dialog box shown in Figure 3.2.

Figure 3.2: A successful conversion of the private key using PuTTYgen.

   c. Click "Save private key," optionally providing a passphrase and specifying a **.ppk** extension (e.g., **ibm_cloud_your.name@ibm.com.ppk**).

2. Launch PuTTY, and specify these settings:
   a. In the Session section, specify the IP address of your newly created instance.
   b. In the Data section, specify **idcuser** as the username.
   c. In the Auth section, select "Allow attempted changes of user name," and specify the path to the private keyfile you saved in Step 2 (**ibm_cloud_your.name@ibm.com.ppk**). Figure 3.3 shows these settings.

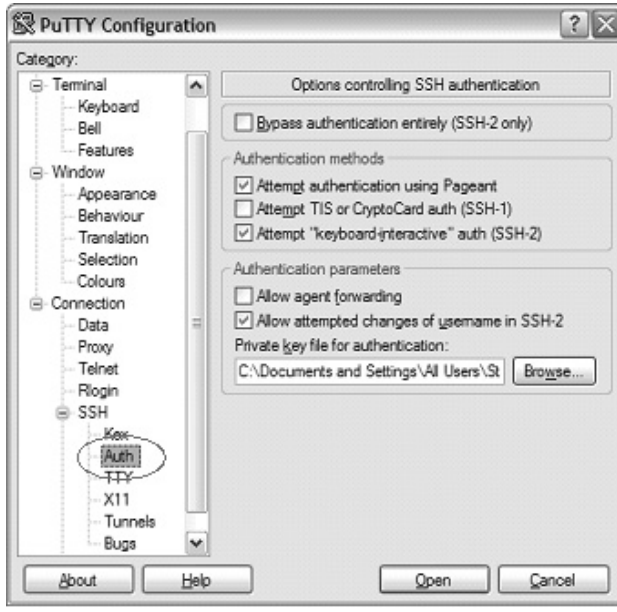

Figure 3.3: SSH authentication settings in PuTTY.

d. In the X11 section, select "Enable X11 forwarding" (Figure 3.4).

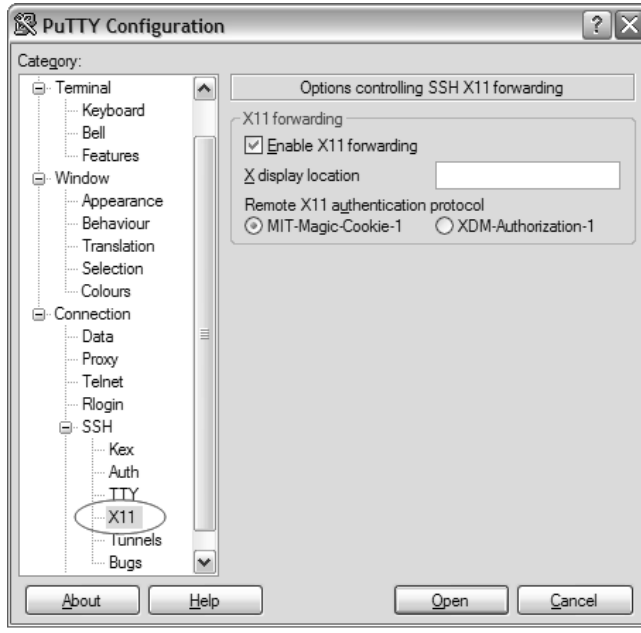

Figure 3.4: X11 forwarding settings in PuTTY.

3. Name and save this session for later use (in the Session section).
4. Click **Open** to securely connect to your instance, which will let you interact with the instance through the command line shell.

Configure WinSCP

To configure WinSCP, follow these steps:

1. Launch WinSCP.
2. In the New Session dialog, specify the IP address of your newly created instance, and enter **idcuser** as the username (Figure 3.5).

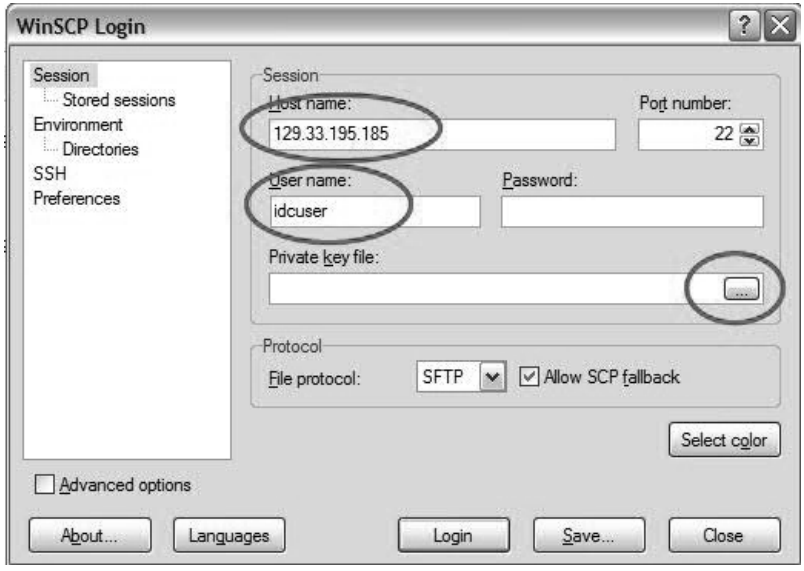

Figure 3.5: Creating a login session in WinSCP.

3. Click the "Private key file" selection button, and choose the **.ppk** file you saved previously (**ibm_cloud_your.name@ibm.com.ppk**). Click **Open**.
4. Click **Login**, and wait for the Login process dialog. When the WinSCP session window appears, you can copy files back and forth.

Test X-Windows on Your Client

To test X-Windows on the client, follow these steps:

1. Launch Xming, and confirm that the Xming icon appears in the system tray.
2. Launch a PuTTY console, and connect to your cloud instance as before.
3. Enter **xterm &** into the command line:

```
$ xterm &
```

4. If your X-Windows client is working, an Xterm window should appear (Figure 3.6).

Figure 3.6: Xterm successfully opened on the local desktop.

Assemble Your Software

If you do not have the required software handy, you can download it in advance from IBM's Passport Advantage®, at *http://www.ibm.com/software/passportadvantage*. You will need your customer Passport Advantage account number to do so.

The sections that follow describe how to achieve either a Cognos 8 Business Intelligence or a Cognos 10 Business Intelligence installation. In most cases, the steps are the same or very similar for these two versions. Illustrative screenshots are provided from both versions of Cognos, so there might be minor differences between what is in this book and what you see on your screen. However, these differences should not be confusing, and any major variations between the two installations are noted explicitly.

In addition, both Cognos 8 and Cognos 10 BI are available in 32-bit and 64-bit versions. Again, the instructions are identical in most cases, and any significant differences are noted explicitly.

For these instructions, you'll need the following software for a Cognos 8 installation:

- IBM Cognos 8 Business Intelligence Server 32-bit 8.4.1 Linux x86 Multilingual
- IBM Cognos 8 Business Intelligence Samples 8.4.1 Multiplatform Multilingual
- WebSphere Application Server V7.0 for Linux on x86Series, 32-bit Support
- WebSphere Application Server V7.0 Supplements Application Client, IBM HTTP Server, Web Server Plug-ins, Installation Factory, Migration Tool, IBM Support Assistant and Update Installer for Linux on x86Series, 32-bit Support Multilingual
- IBM DB2 Express Edition 9.7 for Linux on 32-bit AMD and Intel® systems (x86)
- Any required Fix Packs for the above software

If you are installing Cognos 8 into a 64-bit Red Hat instance, replace the components listed above with the 64-bit equivalent. WebSphere Application Server V6.1 is also supported with Cognos 8 by following the instructions in this chapter, with very little modification.

For a Cognos 10 installation, you'll need the following software:

- IBM Cognos 10 Business Intelligence Server 32-bit 10.0 Linux x86 Multilingual
- IBM Cognos 10 Business Intelligence Samples 10.0 Multiplatform Multilingual
- WebSphere Application Server V7.0 for Linux on x86Series, 32-bit Support
- WebSphere Application Server V7.0 Supplements Application Client, IBM HTTP Server, Web Server Plug-ins, Installation Factory, Migration Tool,

IBM Support Assistant and Update Installer for Linux on x86Series, 32-bit Support Multilingual
- IBM DB2 Express Edition 9.7 for Linux on 32-bit AMD and Intel systems (x86)
- Any required Fix Packs for the above software

If you are installing Cognos 10 into a 64-bit Red Hat instance, replace the preceding components with the 64-bit equivalent.

Using your file transfer software (e.g., WinSCP), create a directory on your cloud instance to store all the assembled software—**/home/idcuser/installs**—and transfer the software to that directory. The software will likely be stored in either tarball format (with the extension **.tar.gz**) or ISO format (extension **.iso**). In either case, you'll want to extract the software into convenient directories for the remainder of the installation process. The instructions that follow assume the software is in the locations listed in Table 3.1.

Table 3.1: Software Locations for the Installation Process

Software	Directory
DB2	/home/idcuser/installs/db2
WebSphere Application Server	/home/idcuser/installs/was
IBM HTTP Server	/home/idcuser/installs/ihs
Cognos 8 or Cognos 10 BI Server	/home/idcuser/installs/cognos
Cognos 8 or Cognos 10 Samples	/home/idcuser/installs/samples

Extract the Software from Tarball Files

If your software was provided as tarball files, connect to your cloud instance using PuTTY. Then, follow these steps from the command line to extract the software from the tarball files into the target directories:

1. Create the directories for the software:

```
$ cd /home/idcuser/installs
$ mkdir db2
$ mkdir was
$ mkdir ihs
$ mkdir cognos
$ mkdir samples
```

2. Extract DB2, substituting the correct name for your DB2 tarball:

```
$ cd /home/idcuser/installs/db2
$ tar zxvf ../DB2_9.7_Linux.tar.gz
```

3. Extract WebSphere Application Server, substituting the correct name for your WebSphere tarball:

```
$ cd /home/idcuser/installs/was
$ tar zxvf ../WAS_7.0_Linux.tar.gz
```

4. Extract IBM HTTP Server, substituting the correct name for your HTTP Server tarball:

```
$ cd /home/idcuser/installs/ihs
$ tar zxvf ../IHS_7.0_Linux.tar.gz
```

5. Extract the Cognos 8 or Cognos 10 Business Intelligence Server software, substituting the correct name for your Cognos tarball:

```
$ cd /home/idcuser/installs/cognos
$ tar zxvf ../Cognos_BI_Server_Linux.tar.gz
```

6. Extract the Cognos 8 or Cognos 10 BI Sample reports and data, substituting the correct name for your Cognos Samples tarball:

```
$ cd /home/idcuser/installs/samples
$ tar zxvf ../Cognos_BI_Samples_Linux.tar.gz
```

Extract the Software from ISO Files

If your software was provided as ISO files, you will need to connect to the cloud instance using PuTTY and, from the command line, “mount” the individual ISO files to a directory (which makes the contents of the files within the ISO file accessible). You can then copy the software to the target directories. Follow these steps:

1. Create the directories for the software:

```
$ cd /home/idcuser/installs
$ mkdir db2
$ mkdir was
$ mkdir ihs
$ mkdir cognos
$ mkdir samples
```

2. Create a directory to be used as the mount point for the ISO files:

```
$ sudo mkdir /mnt/disk
```

3. Copy the DB2 software, substituting the correct name for your DB2 ISO file:

```
$ sudo mount -o loop DB2_9.7_Linux.iso /mnt/disk
$ cp -R /mnt/disk/* /home/idcuser/installs/db2
$ sudo umount -fl /mnt/disk
```

4. Copy the WebSphere Application Server software, substituting the correct name for your WebSphere ISO file:

```
$ sudo mount -o loop WAS_7.0_Linux.iso /mnt/disk
$ cp -R /mnt/disk/* /home/idcuser/installs/was
$ sudo umount -fl /mnt/disk
```

5. Copy the IBM HTTP Server software, substituting the correct name for your HTTP Server ISO file:

```
$ sudo mount -o loop IHS_7.0_Linux.iso /mnt/disk
$ cp -R /mnt/disk/* /home/idcuser/installs/ihs
$ sudo umount -fl /mnt/disk
```

6. Extract the Cognos 8 or Cognos 10 BI software, substituting the correct name for your Cognos ISO file:

```
$ sudo mount -o loop Cognos_BI_Server_Linux.iso /mnt/disk
$ cp -R /mnt/disk/* /home/idcuser/installs/cognos
$ sudo umount -fl /mnt/disk
```

7. Extract the Cognos 8 or Cognos 10 BI Sample reports and data, substituting the correct name for your Cognos Samples tarball:

```
$ sudo mount -o loop Cognos_BI_Samples_Linux.iso /mnt/disk
$ cp -R /mnt/disk/* /home/idcuser/installs/samples
$ sudo umount -fl /mnt/disk
```

Set Up the Database and Web Server

Take the following steps to set up the database and Web server.

Install DB2 Express

We'll use the DB2 Express product as the database for our Cognos installation:

1. Run the DB2 Setup Wizard program, shown in Figure 3.7:

```
$ sudo /home/idcuser/installs/db2/exp/db2setup
```

Figure 3.7: The DB2 Setup Wizard.

2. Select "Install a Product" and then "Install New."
3. From the welcome page, click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Set the installation directory to **/opt/ibm/db2**.
6. On the DAS page, specify the administrator user name and password you want to use. (Our example uses **dasusr1** for both.)
7. On the Instance Setup page, select "Create Instance" and click **Next**.

8. On the Instance Owner page, specify a user name and password. Figure 3.8 shows **db2inst1** entered for both.

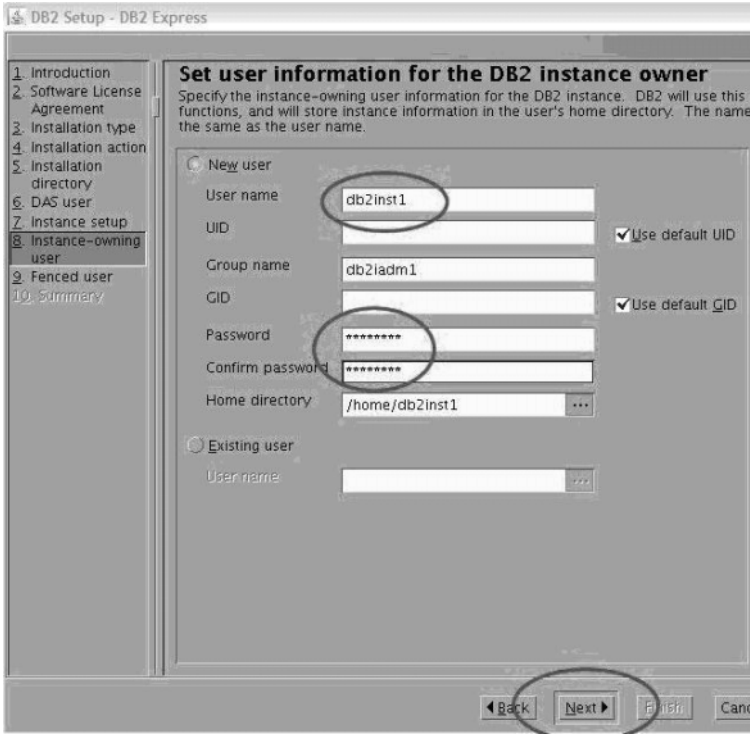

Figure 3.8: Setting up the DB2 instance owner.

9. On the Fenced User page, specify a user name and password. (Our example uses **db2fenc1** for both.)
10. Click **Finish** on the Summary page, and let the automated installation finish. You will see the message shown in Figure 3.9 once setup is complete.

Figure 3.9: A successful DB2 installation.

11. Update your DB2 environment by modifying your **~/.bash_profile** to add the following lines to the end of the file:

```
; Add the following to the bottom:
DB2DIR=/opt/ibm/db2; export DB2DIR
DB2INSTANCE=db2inst1; export DB2INSTANCE
PATH=$PATH:$DB2DIR/bin; export PATH
LD_LIBRARY_PATH=$DB2DIR/lib32; export LD_LIBRARY_PATH
sudo ksh $DB2DIR/cfg/db2profile
```

12. Close and start a new PuTTY session, and validate your DB2 installation:

```
$ sudo db2val
```

If all goes well, you should see output similar to the following:

```
DBI1379I  The db2val command is running. This can take several minutes.
DBI1335I  Installation file validation for the DB2 copy installed at
 /opt/ibm/db2/V9.7_01 was successful.
DBI1343I  The db2val command completed successfully. For details,
 see the log file /tmp/db2val-05_15_01:21:04.log.
```

Install WebSphere Application Server

We'll use WebSphere Application Server as the Web application server for our Cognos BI installation. Before doing so, because the installer will use the browser, ensure that the browser is working and initialized by launching it:

```
$ firefox
```

Next, follow these steps:

1. Run the WebSphere Application Server installation setup wizard:

```
$ sudo /home/idcuser/installs/websphere/launchpad.sh
```

2. Enable the "Enable Administrative Security" option, and specify **idcuser** for the administrative user. Follow the remaining on-screen instructions to complete the installation.
3. Update the **JAVA_HOME** and **PATH** environment variables to reflect the WebSphere JRE and WebSphere bin folder, respectively, by making the following changes to the end of the user's **.bash_profile** file:

```
$ vi /home/idcuser/.bash_profile
# add to the bottom of .bash_profile
# JAVA_HOME=/opt/ibm/websphere/appserver/java; export JAVA_HOME
$ PATH=$JAVA_HOME/jre/bin:$PATH; export PATH
```

4. Exit the PuTTY console, and open a new connection to your cloud instance to have the new environment variables take effect.

Install IBM HTTP Server

Now that WebSphere Application Server is set up as our Web application server, let's use the IBM HTTP Server (IHS) Web server for the Cognos installation.

1. Run the IHS Installation Setup Wizard (Figure 3.10), using the following command:

```
$ sudo /home/idcuser/installs/ihs/launchpad.sh
```

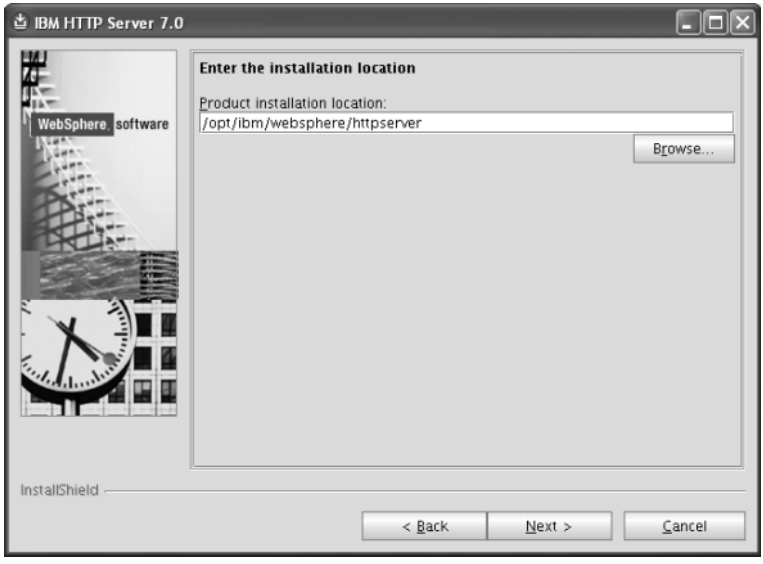

Figure 3.10: The IBM HTTP Server Installation Setup Wizard.

2. Add folder **/usr/lib** to the beginning of the **LD_LIBRARY_PATH** environment variable at the bottom of **.bash_profile**. This will prevent a

"Couldn't load XPCOM" error when trying to launch certain X11-based applications.

```
$ vi ~/.bash_profile
LD_LIBRARY_PATH=/usr/lib:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH
```

3. Point the browser to your machine instance:

```
$ firefox http://localhost &
```

As an alternative, run the browser from your Windows client PC and connect to the public address associated with the machine instance. For example, use "http://*vhost.ibm.com*," replacing *vhost.ibm.com* with the address that appears in the IBM Cloud Control Panel. Doing so should produce the default IBM HTTP Server page (Figure 3.11).

Figure 3.11: The default IBM HTTP Server home page.

Install Required Maintenance Updates

At the time of this writing, the conformance for Cognos BI Server requires WebSphere 7 Fix Pack 5 with JDK6SR5 to be installed. You can install these maintenance packages using the Update Installer program for WebSphere, as follows:

1. On your cloud instance, create a folder to contain the required files:

```
$ mkdir -p /home/idcuser/installs/websphere/fixpack11/
```

2. Download the Update Installer and the Fix Pack files listed in Table 3.2, and transfer them to the **/home/idcuser/installs** directory.

Table 3.2: Update Installer and Fix Pack Files

WebSphere 7 Update Installer	ftp://public.dhe.ibm.com/software/websphere/appserv/support/tools/UpdateInstaller/7.0.x/LinuxIA32/7.0.0.11-WS-UPDI-LinuxIA32.tar.gz
WebSphere 7 Fix Pack 11	ftp://ftp.software.ibm.com/software/websphere/appserv/support/fixpacks/was70/cumulative/cf70011/LinuxX32/7.0.0-WS-WAS-LinuxX32-FP0000011.pak
JDK6SR7	ftp://ftp.software.ibm.com/software/websphere/appserv/support/fixpacks/was70/cumulative/cf70011/LinuxX32/7.0.0-WS-WASSDK-LinuxX32-FP0000011.pak

3. Extract the Update Installer software to **home/idcuser/installs/websphere/fixpack11**:

```
$ cd /home/idcuser/installs/websphere/fixpack11/
$ tar zxvf 7.0.0.11-WS-UPDI-LinuxIA32.tar.gz
```

4. Install Update Installer for WebSphere using the installation wizard, shown in Figure 3.12:

```
$ sudo /home/idcuser/installs/websphere/fixpack11/UpdateInstaller/
  install &
```

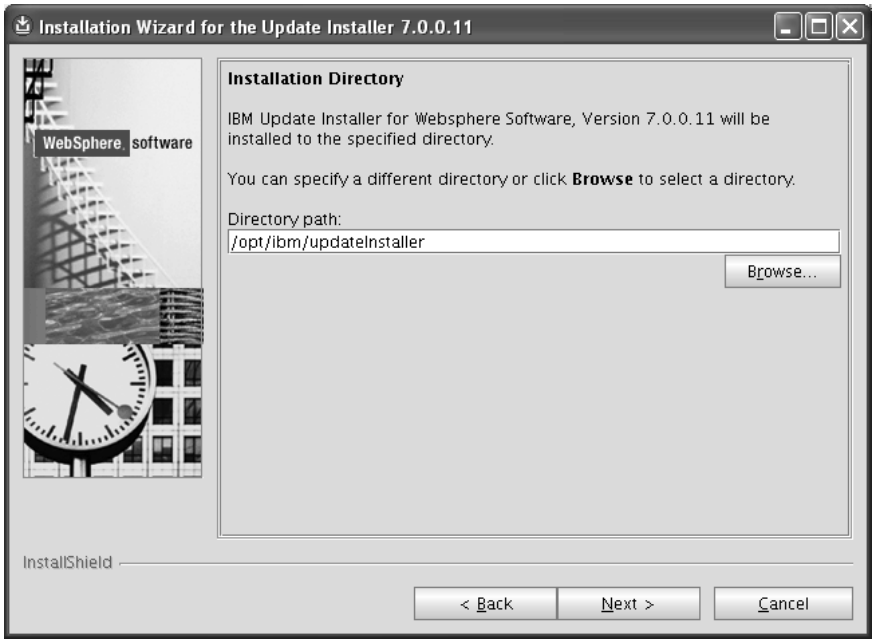

Figure 3.12: The WebSphere Update Installer Installation Wizard.

5. Copy the Fix Pack files to the Update Installer's "maintenance" location, and then run the Update Installer:

```
$ sudo cp /home/idcuser/installs/websphere/fixpack11/7.0.0-WS-WAS-
  LinuxX32-FP0000011.pak /opt/ibm/updateInstaller/maintenance/
$ sudo cp /home/idcuser/installs/websphere/fixpack11/7.0.0-WS-WASSDK-
  LinuxX32-FP0000011.pak /opt/ibm/updateInstaller/maintenance/
$ sudo /opt/ibm/updateInstaller/update.sh &
```

The installer will go through a guided process to detect all installed WebSphere components, including Application Server, HTTP Server, and HTTP Plug-ins, and to apply the Fix Packs. Depending on the components that need to be updated, you might need to run Update Installer multiple times to ensure all Fix Packs are applied.

Configure IBM HTTP Server

A simple Cognos BI Server has two directories that need to be mapped in your Web server: **cgi-bin** and **webcontent**.

1. Edit the **/opt/ibm/websphere/httpserver/conf/httpd.conf** file, and append the following to the existing file contents:

```
ScriptAlias /ibm/cognos/cgi-bin "/opt/ibm/cognos/cgi-bin/"
<Directory "/opt/ibm/cognos/cgi-bin">
  Options FollowSymLinks ExecCGI
  AllowOverride All
  Order allow,deny
  Allow from all
</Directory>
Alias /ibm/cognos "/opt/ibm/cognos/webcontent"
<Directory "/opt/ibm/cognos/webcontent">
  Options FollowSymLinks Indexes MultiViews
  AllowOverride None
  Order allow,deny
  Allow from all
</Directory>
```

2. Restart the Apache Web server to pick up the new settings:

```
$ sudo /opt/ibm/websphere/httpserver/bin/apachectl stop
$ sudo /opt/ibm/websphere/httpserver/bin/apachectl start
```

3. Confirm that the directory mapping was set up by starting a local browser and pointing it to your machine instance:

```
$ firefox http://localhost/ibm/cognos &
```

You can also launch a browser on your Windows machine and point it to the IP address of your machine instance. Use the format "http://*vhost.ibm.com*/ibm/cognos," substituting *vhost.ibm.com* with the correct address from the IBM Cloud Control Panel.

You should see the Cognos splash screen, followed by the following error delivered to your browser: "The IBM Cognos gateway is unable to connect to the IBM Cognos BI Server. The server may be unavailable or the gateway may not be correctly configured." This behavior is correct because you have not yet completed the setup of Cognos. It confirms that your virtual mappings are complete.

Set Up Cognos 8 or Cognos 10 BI Server

In this section, you will complete your single-instance deployment by installing and configuring Cognos 8 or Cognos 10 Business Intelligence Server and the standard samples into your cloud instance.

Install Cognos BI Server

1. Create a target folder for your Cognos installation as the root user, and then change its ownership to **idcuser**:

```
$ sudo mkdir –p /opt/ibm/cognos
$ sudo chown idcuser /opt/ibm/cognos
$ sudo chgrp idcuser /opt/ibm/cognos
```

2. Run the Cognos BI installation program (replacing **linuxi38632** with **linuxi38664** if installing the 64-bit version instead of the 32-bit version). This will initiate the standard Installation Wizard process:

```
$ /home/idcuser/installs/cognos/linuxi38632/issetup &
```

3. On the Welcome screen (Figure 3.13), select your language of choice, and click **Next**.

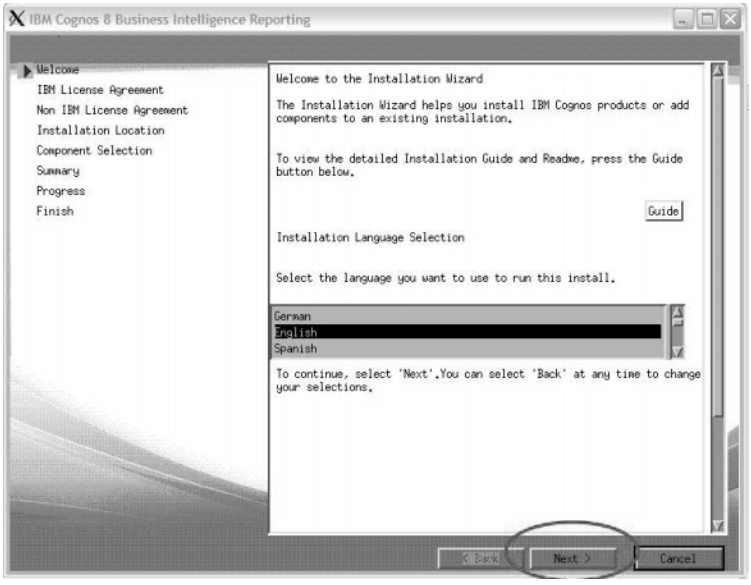

Figure 3.13: The Cognos Installation Wizard.

4. Read and accept the license agreement. Click **Next**.
5. Accept the standard installation location of **/opt/ibm/cognos** (Figure 3.14). Click **Next**. If the installer asks to create a new directory, click **Yes**.

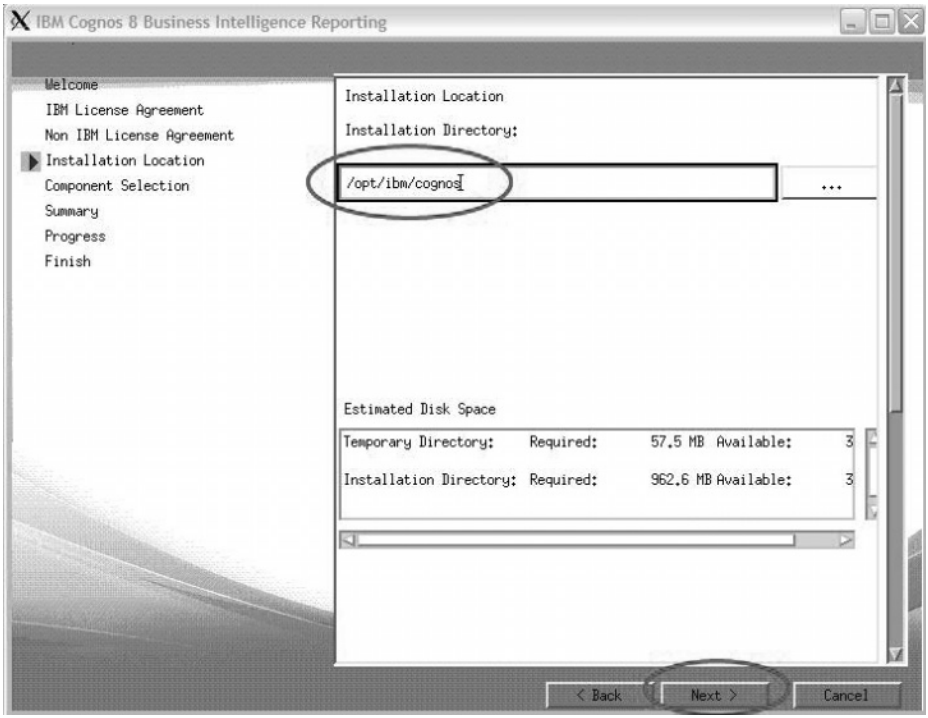

Figure 3.14: Specifying an installation location.

6. On the Component Selection page (Figure 3.15), select "Application Tier Components," "Gateway," and "Content Manager." Click **Next**.

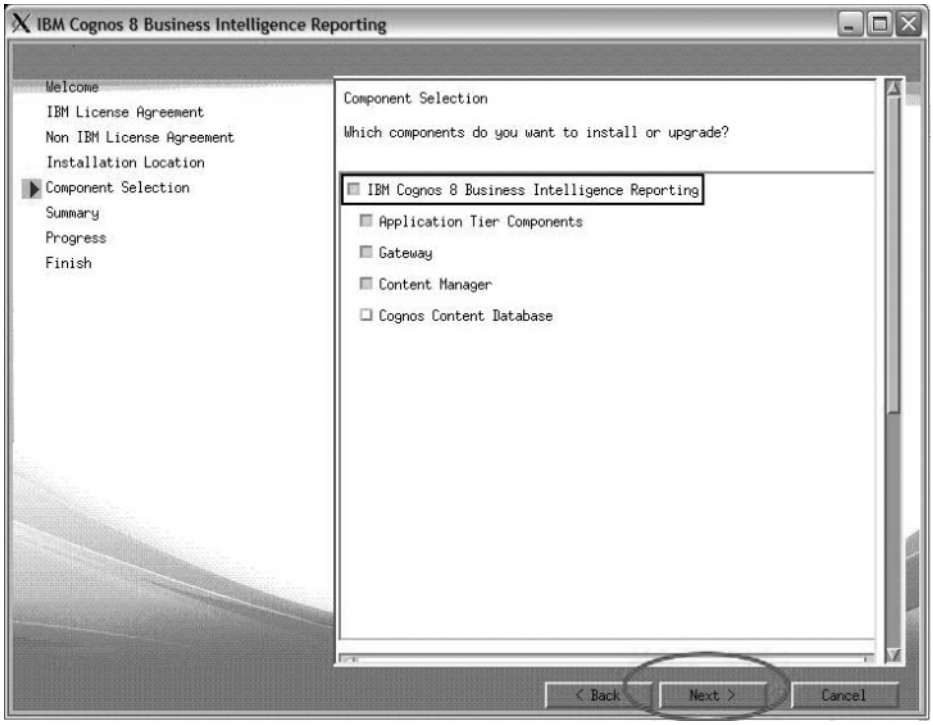

Figure 3.15: Selecting the Cognos components to install.

7. Check the installation summary, and when satisfied click **Next**.

8. Wait for the installation to complete, and then review your installation results (Figure 3.16). Click **Finish**.

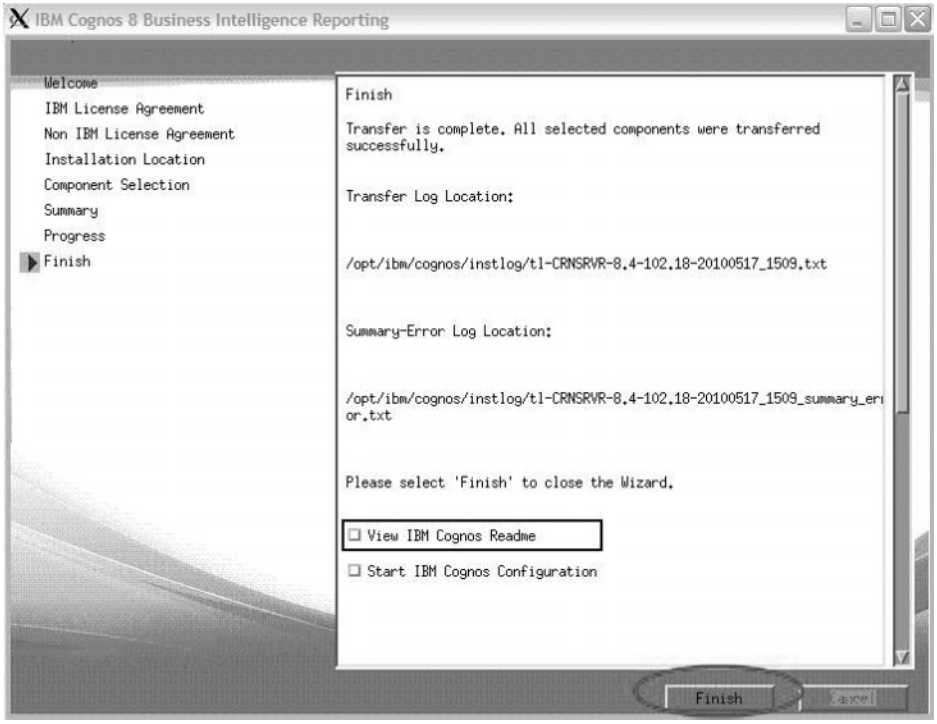

Figure 3.16: A successful Cognos BI installation.

Create the Content Store

Next, create the Cognos content store in the previously created DB2 instance.

1. Add the **idcuser** and **root** users to the **db2adm1** group:

```
$ sudo vi /etc/group
db2iadm1:x:105:root,idcuser
```

   This step is required for Cognos BI startup. It permits both users to directly access various DB2 administration tools.
2. Create a new DB2 command script called **cognos_init_recreateContentStore.sql**:

```
$ vi /home/idcuser/cognos_init_recreateContentStore.sql
```

This script will contain DB2 commands to be executed via command line in step 4. You will also use this script in the "Prepare Image Configuration Scripts" section toward the end of this chapter.

3. Add these DB2 commands to the script you created in the previous step:

```
FORCE APPLICATION ALL
@
DB2STOP
@
DB2START
@
DROP DATABASE "COGNOS"
@
CREATE DATABASE "COGNOS" ALIAS COGNOS
@
CHANGE DATABASE "COGNOS" COMMENT WITH 'Cognos Content Store'
@
CONNECT TO "COGNOS"
@
UPDATE DATABASE CONFIGURATION USING APPLHEAPSZ 1024 DEFERRED
@
UPDATE DATABASE CONFIGURATION USING LOCKTIMEOUT 240 DEFERRED
@
CONNECT RESET
@
CONNECT TO "COGNOS"
@
CREATE BUFFERPOOL "COGNOS_08KBP" IMMEDIATE SIZE 1000 PAGESIZE 8K
@
CREATE BUFFERPOOL "COGNOS_32KBP" IMMEDIATE SIZE 1000 PAGESIZE 32K
@
CONNECT RESET
@
CONNECT TO "COGNOS"
```

```
@
CREATE SYSTEM TEMPORARY TABLESPACE “TSN_SYS_COGNOS” IN DATABASE
   PARTITION GROUP IBMTEMPGROUP PAGESIZE 32K MANAGED BY SYSTEM USING
   (‘CNT_SYS_COGNOS’) BUFFERPOOL “COGNOS_32KBP”
@
CREATE USER TEMPORARY TABLESPACE “TSN_USR_COGNOS” IN DATABASE
   PARTITION GROUP IBMDEFAULTGROUP PAGESIZE 8K MANAGED BY SYSTEM
   USING (‘CNT_USER_COGNOS’) BUFFERPOOL “COGNOS_08KBP”
@
CREATE REGULAR TABLESPACE “TSN_REG_COGNOS” IN DATABASE PARTITION
   GROUP IBMDEFAULTGROUP PAGESIZE 8K MANAGED BY SYSTEM USING (‘CNT_
   REG_COGNOS’) BUFFERPOOL “COGNOS_08KBP”
@
ALTER TABLESPACE “TSN_REG_COGNOS” DROPPED TABLE RECOVERY ON
@
CONNECT RESET
@
CONNECT TO “COGNOS”
@
CREATE SCHEMA db2COGNOS AUTHORIZATION db2inst1
@
COMMENT ON SCHEMA db2COGNOS IS ‘Cognos Content Store’
@
GRANT CREATETAB,BINDADD,CONNECT,IMPLICIT_SCHEMA ON DATABASE  TO USER
   db2inst1
@
GRANT CREATEIN,DROPIN,ALTERIN ON SCHEMA DB2COGNOS TO USER db2inst1
   WITH GRANT OPTION
@
GRANT USE OF TABLESPACE “TSN_USR_COGNOS” TO db2inst1
@
CONNECT RESET
@
DB2STOP
@
```

4. Run the script from the command line to create the Cognos content store on the DB2 instance:

```
$ db2 -td@ -f ~/cognos_init_recreateContentStore.sql -z
   ~/cognos_createContentStore.log
```

This initializes the DB2 database required for Cognos BI. You will now confirm the successful creation of the DB2 database using the DB2 Control Center.

5. Run the DB2 Control Center:

```
$ db2cc &
```

6. On the DB2 Control Center's View menu, select "Advanced" and click **OK**.
7. Expand "All Databases," as shown in Figure 3.17, and select "COGNOS."

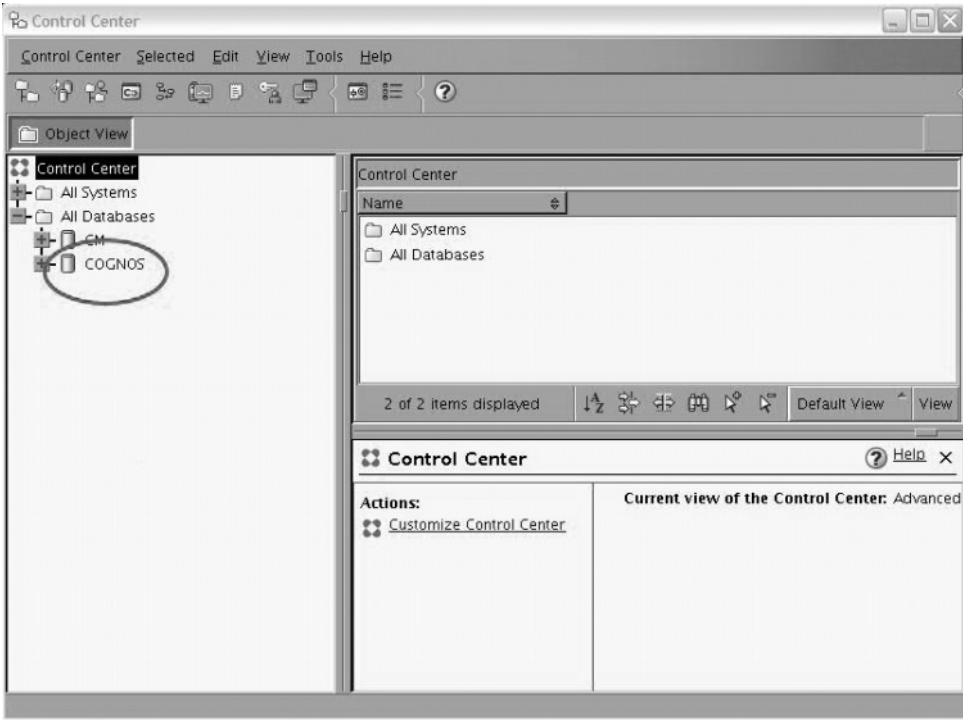

Figure 3.17: The Cognos content store database.

8. Click **Query**, and confirm the database connection. Enter any valid SQL. For example, enter **select count(*) from syscat.tables**, as shown in Figure 3.18. Click **Execute** to confirm the SQL execution results. Close the Control Center.

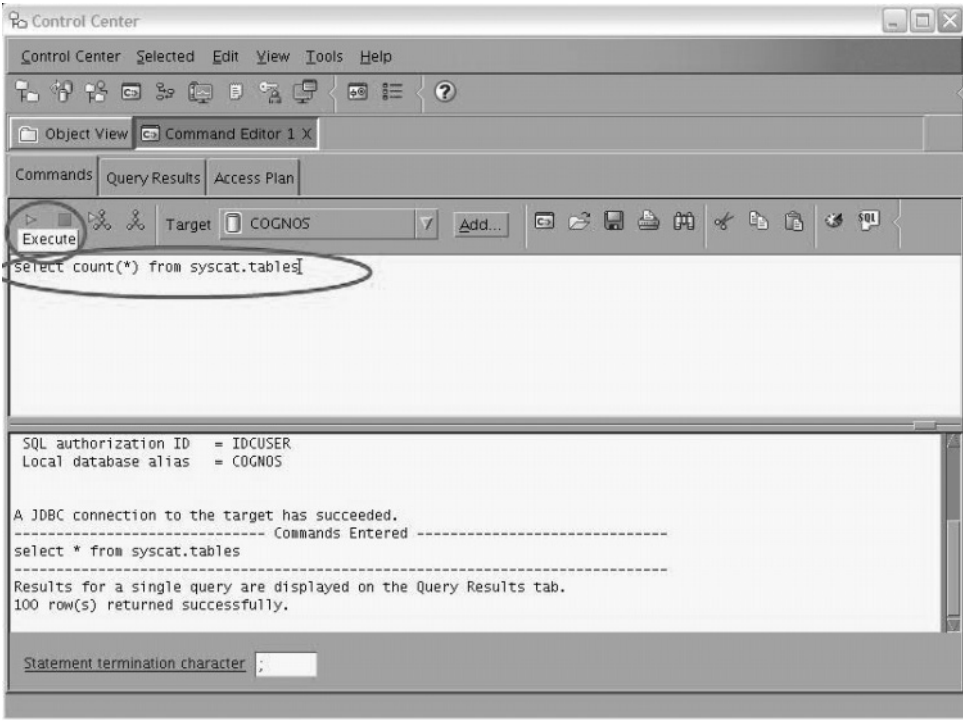

Figure 3.18: Testing the Cognos content store database.

Install IBM JCE Unlimited Strength Jurisdiction Policy Files

By default, the Java™ Cryptography Extension (JCE) is shipped with restricted or limited ciphers. To use advanced cryptography, you must apply unlimited jurisdiction policy files. This step can be omitted for 32-bit installs; however, updating the JCE is required for 64-bit installations, to avoid cryptographic keysize errors when keys are generated within Cognos Configuration. Before downloading these policy files, back up the existing policy files (**local_policy.jar** and **US_export_policy.jar** in directory **jre/lib/security/**), in case you want to restore the original files later.

You can obtain unlimited jurisdiction policy files by completing the following steps:

1. Go to the Web site *http://www.ibm.com/developerworks/java/jdk/security/index.html*.
2. Click **Java 1.4.2** (or the latest 1.4.*x* version available on the page).
3. Click **IBM SDK Policy files** to display the "Unrestricted JCE Policy files for SDK 1.4" Web site.
4. Enter your user ID and password, or register with IBM to download the policy files. The policy files, **local_policy.jar** and **US_export_policy.jar**, will be downloaded onto your machine.
5. Create a folder on your cloud instance to store the files, and then copy the extracted **.jar** files to it using WinSCP:

```
$ mkdir /home/idcuser/installs/jce
```

6. As the root user, copy the **.jar** files to the actual **$JAVA_HOME/jre/lib/security** directory:

```
$ sudo cp -rf $JAVA_HOME/jre/lib/security $JAVA_HOME/jre/lib/
  security_bak
$ sudo cp /home/idcuser/installs/jce/*.jar $JAVA_HOME/jre/lib/
  security
```

You have now created the environment that lets you proceed with configuring Cognos BI!

Configure Cognos BI Server

You are ready to configure the Cognos 8 or Cognos 10 BI Server to use the DB2 content store just created. To do so, follow these steps:

1. Add the **LD_PRELOAD** symbol to the bottom of **.bash_profile**:

```
LD_PRELOAD=/usr/lib/libfreetype.so.6; export LD_PRELOAD
```

For a 64-bit install, use the following **LD_PRELOAD** setting instead:

```
LD_PRELOAD=/usr/lib64/libfreetype.so.6; export LD_PRELOAD
```

2. If the **LD_PRELOAD** symbol is not set, you might encounter the error "java.lang.UnsatisfiedLinkError: /usr/lib/libXft.so.2: undefined symbol: FT_GlyphSlot_Embolden error."
3. Back up the Cognos configuration folder:

```
$ cp -rf /opt/ibm/cognos/configuration /opt/ibm/cognos/configuration_bak
```

4. Copy the DB2 JDBC drivers from the DB2 library directory to the Cognos library directory:

```
$ sudo cp $DB2DIR/java/db2jcc.jar
  /opt/ibm/cognos/webapps/p2pd/WEB-INF/lib
$ sudo cp $DB2DIR/java/db2jcc_license_cu.jar
  /opt/ibm/cognos/webapps/p2pd/WEB-INF/lib
```

5. Copy the Bouncy Castle security keys provider **.jar** file from the Cognos directory to the **JAVA_HOME** directory used by WebSphere (for Cognos 8, replace JRE version "6.0" with "1.5.0"):

```
$ sudo cp /opt/ibm/cognos/bin/jre/6.0/lib/ext/bcprov-jdk14-134.jar
  $JAVA_HOME/jre/lib/ext/
```

6. Launch the Cognos Configuration tool to start the configuration process:

```
$ /opt/ibm/cognos/bin/cogconfig.sh &
```

For a 64-bit install, use the following line instead:

```
$ opt/ibm/cognos/bin64/cogconfig.sh &
```

7. Under Environment > IBM Cognos Services, right-click the default "Cognos" service and then click **Delete** (Figure 3.19).

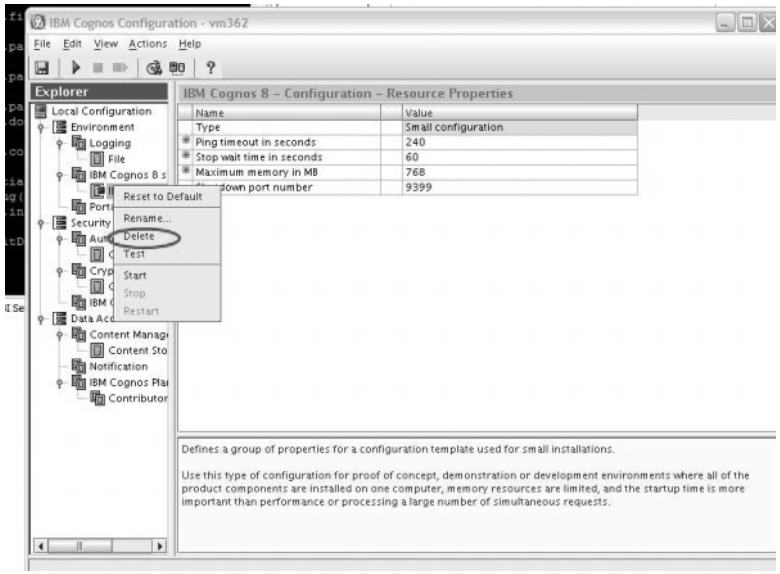

Figure 3.19: Deleting the default Cognos content store.

8. Under Data Access > Content Manager > Content Store, set the database to "localhost:50000" and the database name to "COGNOS" (Figure 3.20). Click the box beside the "User ID and password" field to specify your DB2 username and password (in this example, **db2inst1** for both).

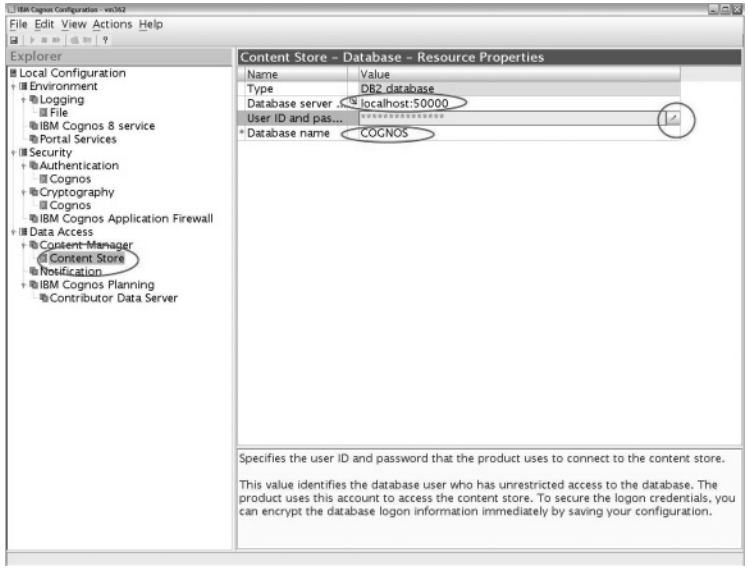

Figure 3.20: Connecting to the DB2 content store.

9. Test your database settings by right-clicking the Content Store and clicking **Test**. If all goes well, you should see the successful-test screen shown in Figure 3.21. Click **Close**.

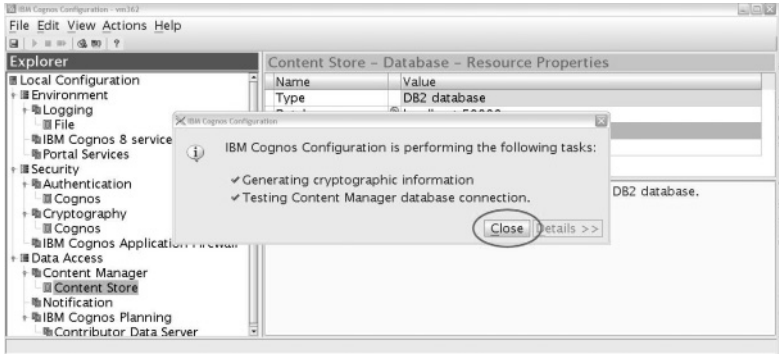

Figure 3.21: Testing the content store.

10. Go to Local Configuration > Environment, and verify that all the URLs listed there are correct (Figure 3.22). In particular, make sure the port number is right. (Note that if you are installing under WebSphere Application Server, the default dispatcher port number is 9080 rather than 9300, which is the Tomcat default.)

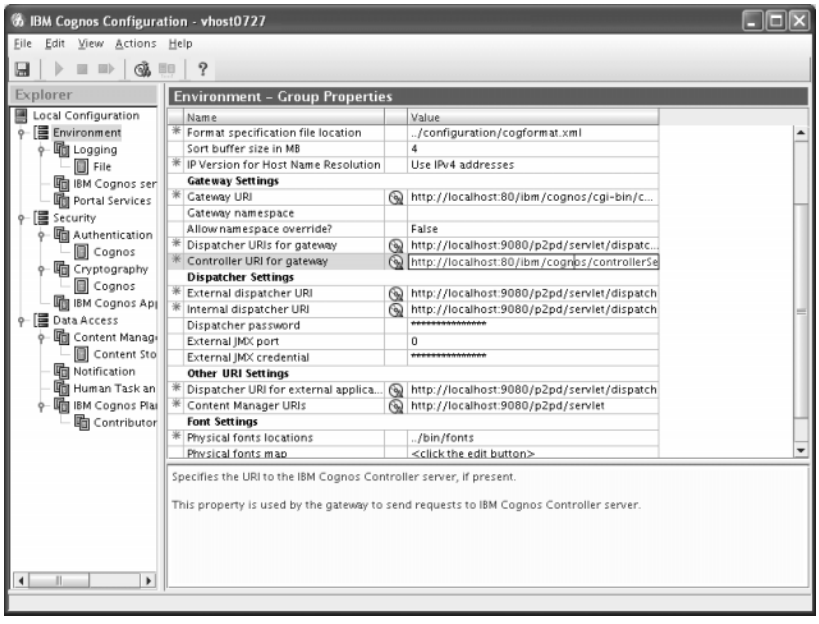

Figure 3.22: Setting the environment URIs.

11. Click **Save** to save your new configuration.
12. From the Actions menu, select "Build Application Files" to launch the Build Application Files Wizard. Select "IBM Cognos," as shown in Figure 3.23. Click **Next**.

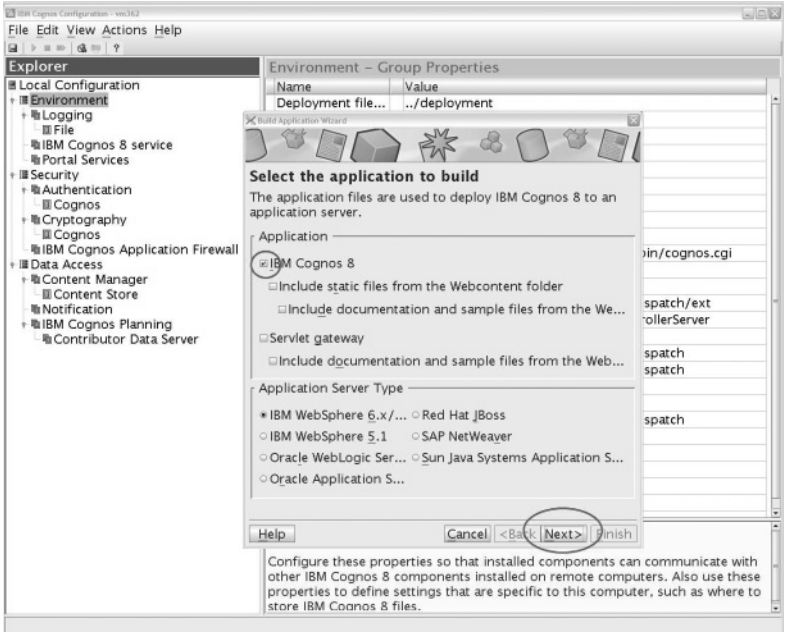

Figure 3.23: Building the Cognos application for WebSphere Application Server.

13. Because you are building for WebSphere Application Server in this example, select "EAR file" for the file type (Figure 3.24). Click **Next**.

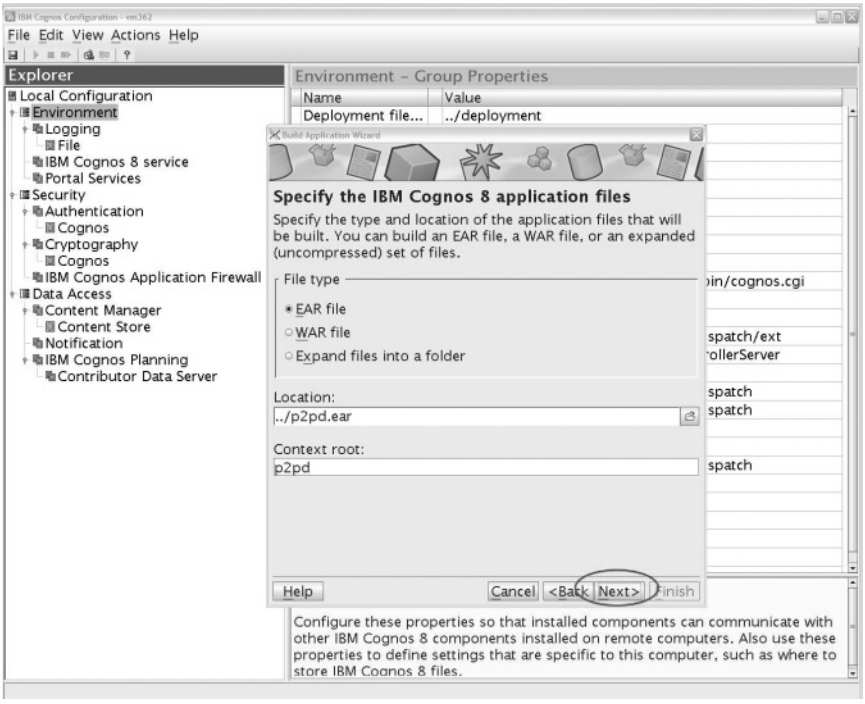

Figure 3.24: Specifying the options for an EAR file.

14. Let the Build Application Files process complete, and then click **Finish**.
15. Exit the Cognos Configuration tool by selecting "Exit" from the File menu.
16. Verify that the Cognos application EAR file, **p2pd.ear**, has been successfully created:

```
$ ls /opt/ibm/cognos/p2pd.ear
```

Install Cognos BI into the WebSphere Application Server

You are now ready to start up your fully configured installation of Cognos BI. Follow these steps:

1. Ensure that the correct permissions are in place:

```
$ sudo chmod -Rf g+wx /opt/ibm/websphere/appserver
```

2. Start WebSphere Application Server:

```
$ /opt/ibm/websphere/appserver/bin/startServer.sh server1 -username
  idcuser -password idcuser
```

3. Launch the WebSphere console to install the Cognos BI application:

```
$ firefox http://localhost:9060/ibm/console &
```

4. If necessary, confirm your browser's security exception. In Firefox, for example, select "Or you can add an exception," then "Add Exception," then "Get Certificate," and finally "Confirm security exception."
5. Enter **idcuser** for the username and password to log into the WebSphere Integrated Solutions Console.
6. Navigate to the Servers > Application Servers tab. Click on your server, as shown in Figure 3.25.

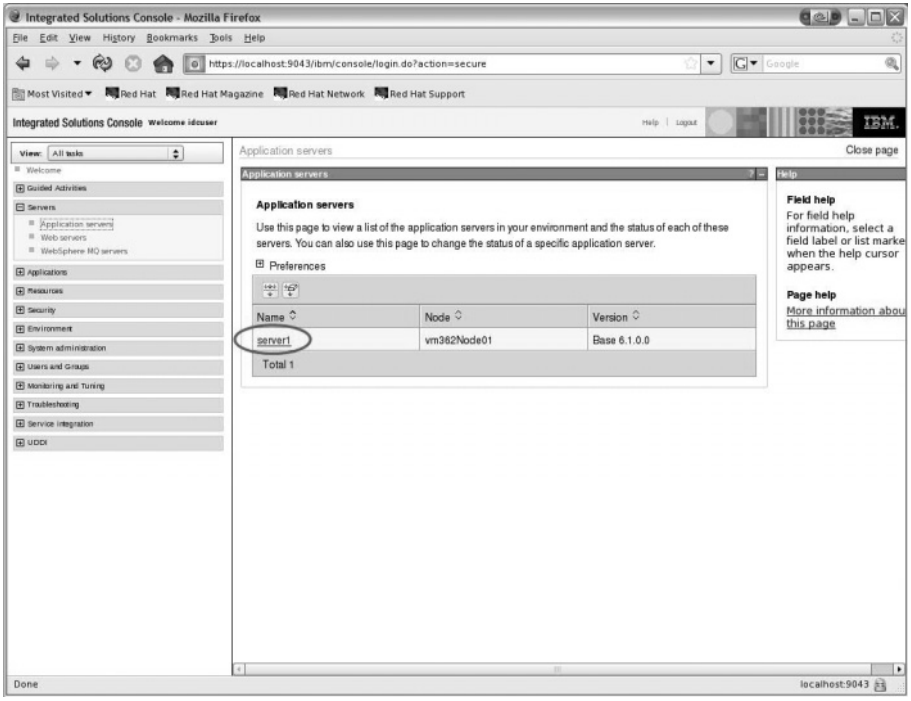

Figure 3.25: The application server in WAS.

7. Under Java and Process Management > Process Definition > Java Virtual Machine, enter 0 (zero) for “Initial Heap Size” and 768 for “Maximum Heap Size” (Figure 3.26). Click **OK** and then **Save**. (This step might not be required for large machine instances.)

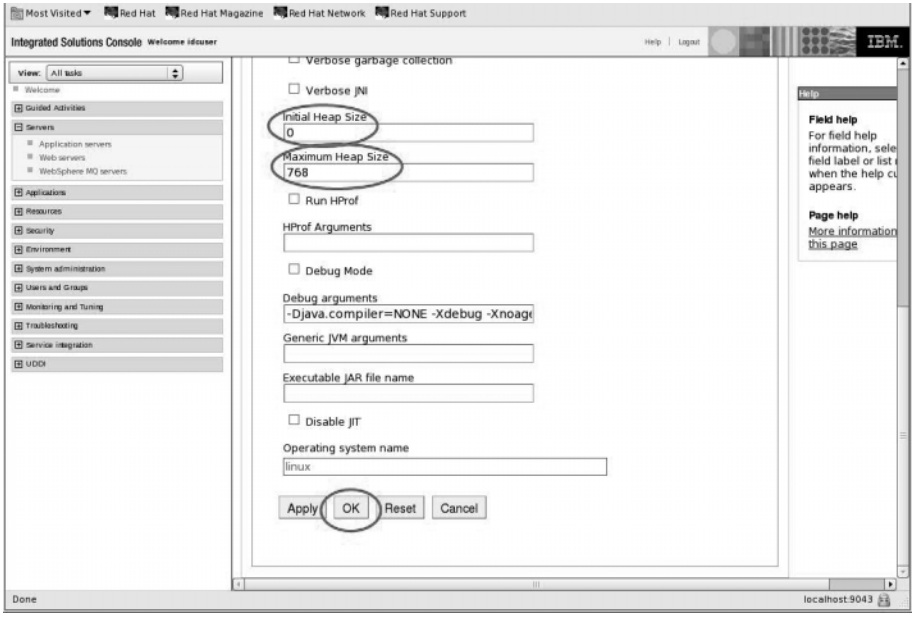

Figure 3.26: Verifying the JVM settings for the WAS application.

8. Under Applications, select “Install New Application” (Figure 3.27) For the “Local File System” path, click **Browse** and select the previously generated EAR file, **/opt/ibm/cognos/p2pd.ear**. Click **Next**.

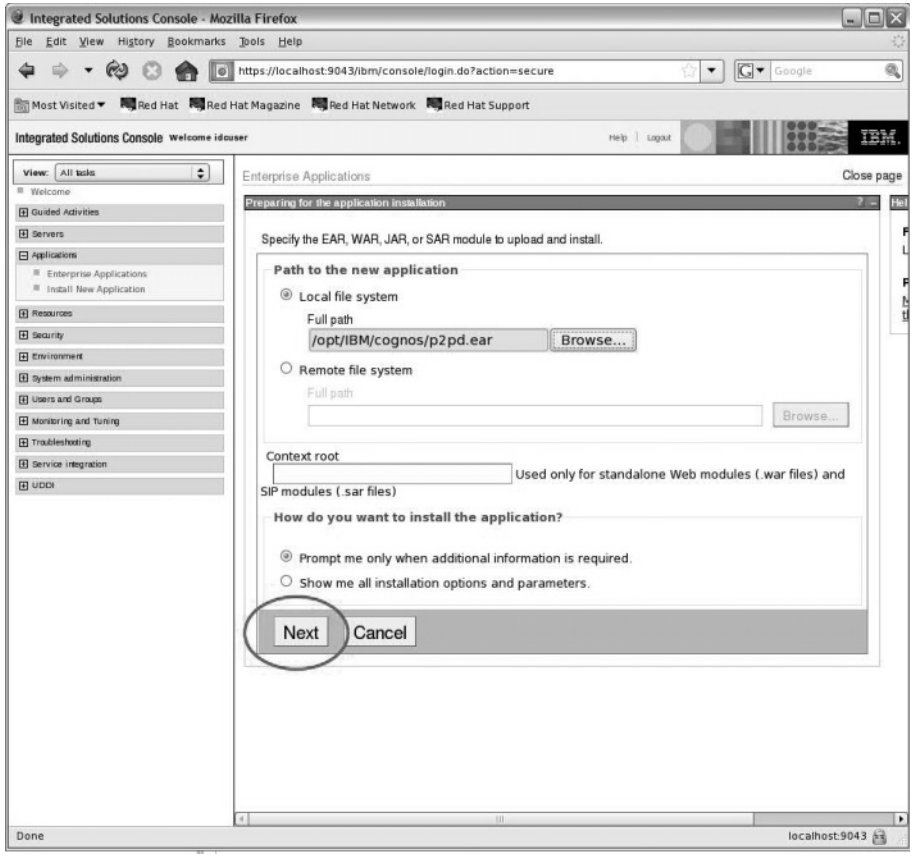

Figure 3.27: Deploying the Cognos EAR file.

9. On the Modules page, select “IBM Cognos” and click **Next**.
10. On the Hosts page, select “IBM Cognos” and click **Next**.
11. At the bottom of the summary page, click **Finish** to start the installation of Cognos into WebSphere Application Server.

12. After the installation is complete, click **Save** to save the changes (Figure 3.28). Then, exit the WAS console.

Figure 3.28: A successful deployment of the Cognos EAR file.

13. In the PuTTY console, start the IBM Cognos WebSphere application. The **LD_LIBRARY_PATH** needs to first be set, so that the required Cognos and DB2 runtime libraries can be located and startup errors can be avoided. Set this variable temporarily, and then restore its original to avoid potential conflicts with other software:

```
$ LD_LIBRARY_PATH_BAK=$LD_LIBRARY_PATH; export LD_LIBRARY_PATH
$ LD_LIBRARY_PATH=/opt/ibm/cognos/bin:$LD_LIBRARY_PATH; export
  LD_LIBRARY_PATH
$ /opt/ibm/websphere/appserver/bin/startServer.sh server1 -username
  idcuser -password idcuser
$ LD_LIBRARY_PATH=$LD_LIBRARY_PATH_BAK; export LD_LIBRARY_PATH
```

For 64-bit support, replace the second line in the preceding statements with this one:

```
LD_LIBRARY_PATH=/opt/ibm/cognos/bin64:$LD_LIBRARY_PATH; export
  LD_LIBRARY_PATH
```

Expect to see the messages shown below upon successful startup:

```
ADMU0116I: Tool information is being logged in file
/opt/ibm/websphere/appserver/profiles/AppSrv01/logs/server1/
  startServer.log
ADMU0128I: Starting tool with the AppSrv01 profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 12357
```

Check **startServer.log** and **/opt/ibm/cognos/cogserver.log** to ensure that Cognos BI started up successfully. You should see the message shown in Figure 3.29, announcing the successful starting of IBM Cognos 8 or 10.

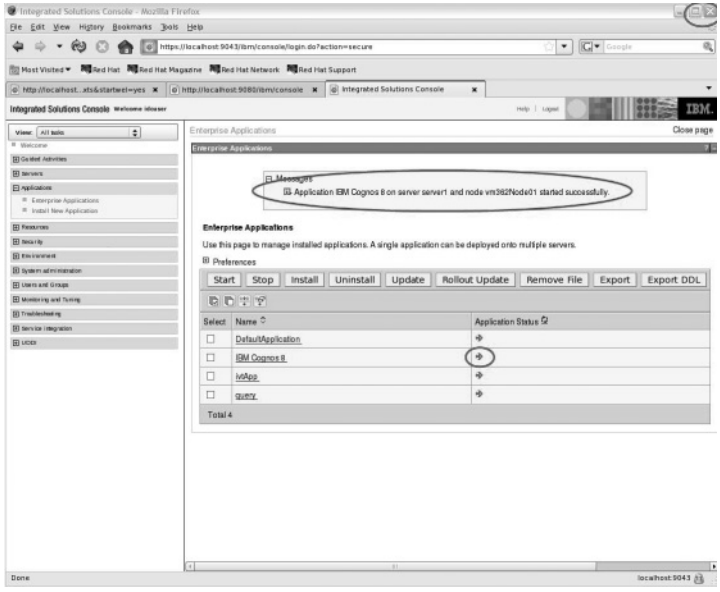

Figure 3.29: Successful startup of the Cognos BI application.

14. Cognos BI should now be running. You can verify this by checking the **stdout** log files generated by the server:

```
$ view /opt/ibm/websphere/appserver/profiles/AppSrv01/logs/server1/
  SystemOut.log
```

You can now point your local browser to the default URL to see the Cognos Connection welcome screen (Figure 3.30):

```
$ firefox http://localhost/ibm/cognos  &
```

As an alternative, use your Windows browser and point it to the IP address or hostname of your machine instance (as displayed via the IBM Cloud Control Panel). For example, use "http://*vhostname.ibm.com*/ibm/cognos," where *vhostname.ibm.com* is your cloud instance's host address.

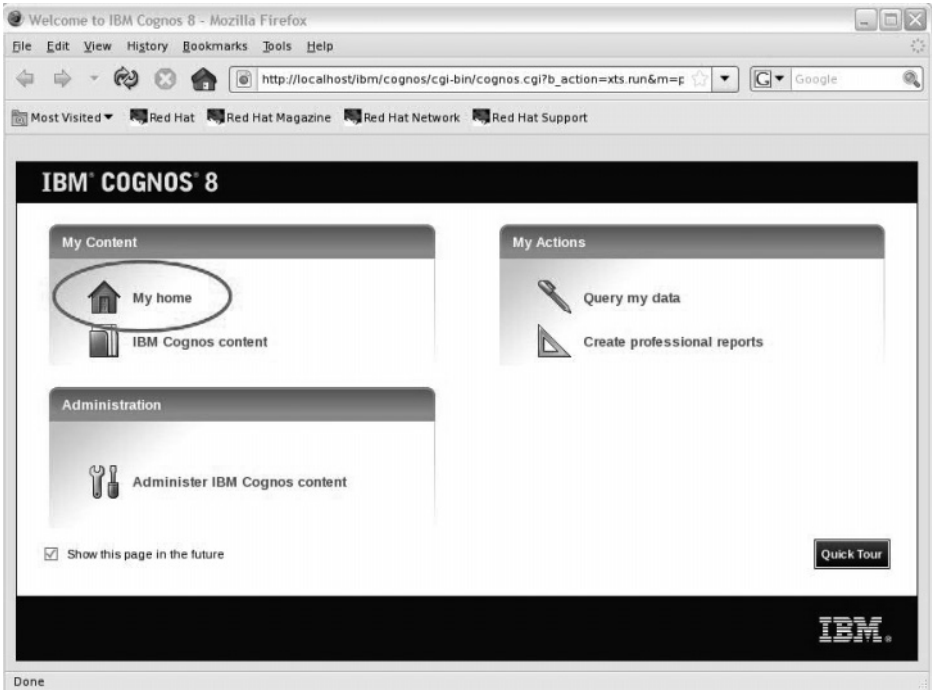

Figure 3.30: The Cognos 8 Business Intelligence default portal page.

Congratulations! You have successfully installed the IBM Cognos 8 or Cognos 10 Business Intelligence Server software to your cloud instance.

Install and Build the Samples Database

A new installation of Cognos BI does not include any sample data sources or reports. In this section, we describe how to install the sample DB2 data sources and Cognos BI reports that are available with the product.

1. Run the Samples installer, replacing **linux38632** with **linux38664** for 64-bit installs:

```
$ cd /home/idcuser/installs/samples/linux38632
$ ./issetup &
```

2. Follow the standard steps to install the samples, ensuring you install them to the same directory where you installed the Cognos software (**/opt/ibm/cognos**, as shown in Figure 3.31).

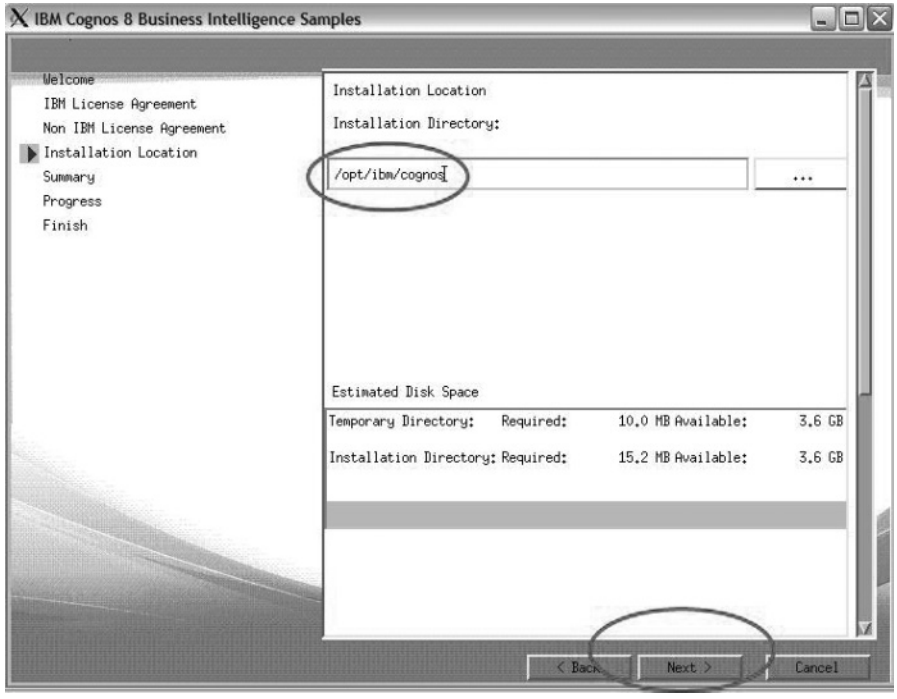

Figure 3.31: The Cognos Samples Installation Wizard.

3. Update the **DB2CODEPAGE** environment variable by editing your **.bash_profile** file:

```
$ vi ~/.bash_profile
```

Add the following line to the file:

```
DB2CODEPAGE=1208; export DB2CODEPAGE
```

4. Create the GS_DB2 database using the **db2cc** utility:

```
$ db2cc &
```

5. From the Control Center view, select "Advanced." Click **OK**.
6. Right-click **All Databases**, and select Create Database > Standard (Figure 3.32).

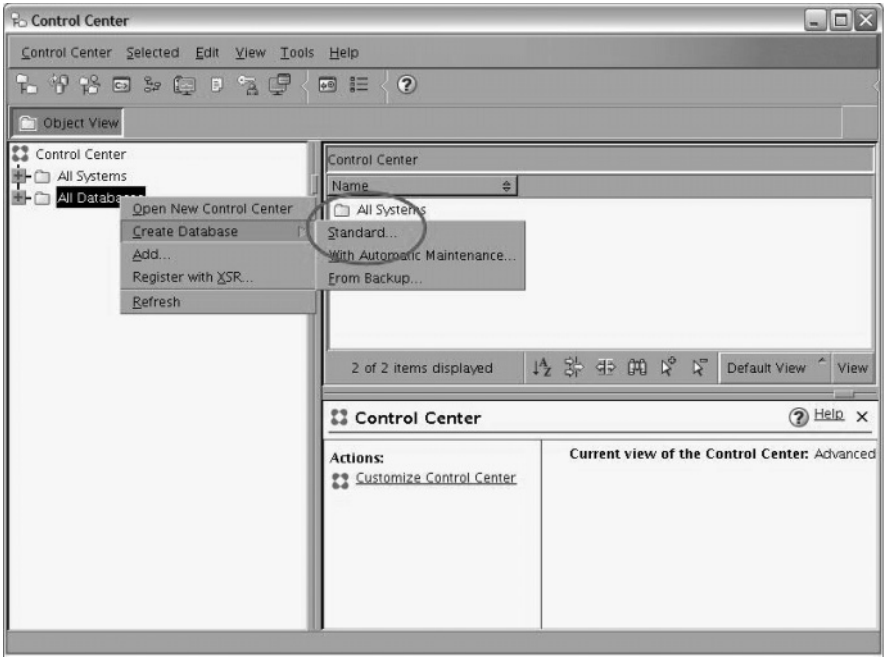

Figure 3.32: Creating a new database for the samples.

7. For the database name, enter **GS_DB** as specified on the Cognos BI Install Guide, and ensure the buffer pool is set to 16K (Figure 3.33). Click **Next**.

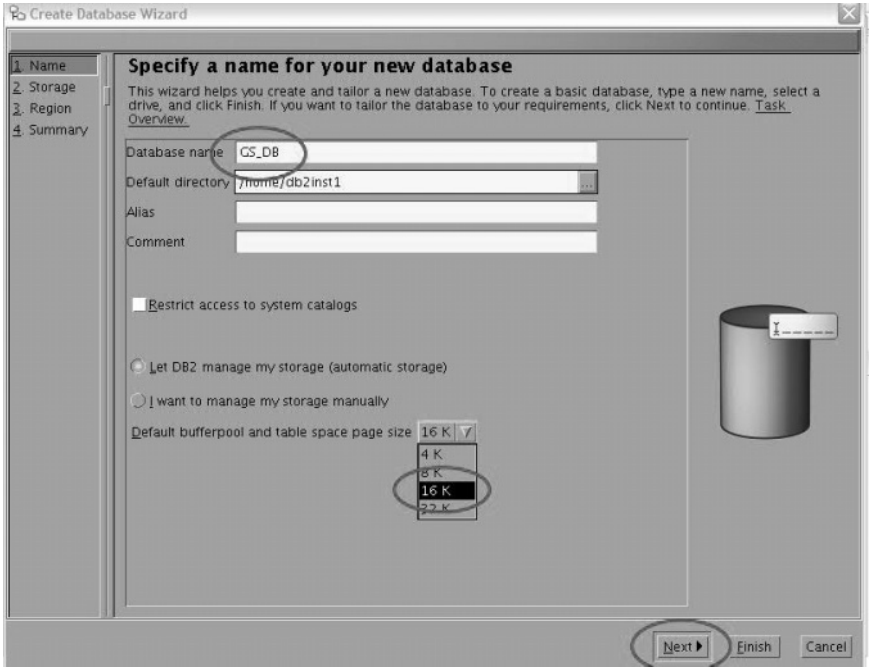

Figure 3.33: Settings for the DB2 samples database.

8. Accept the default storage settings, and click **Next**. Then accept the default region settings, and click **Next**.
9. Let the creation of the database proceed, and close the summary dialog when it is complete.

10. You should now see your new **GS_DB** database under All Databases, as shown in Figure 3.34. Exit the Control Center.

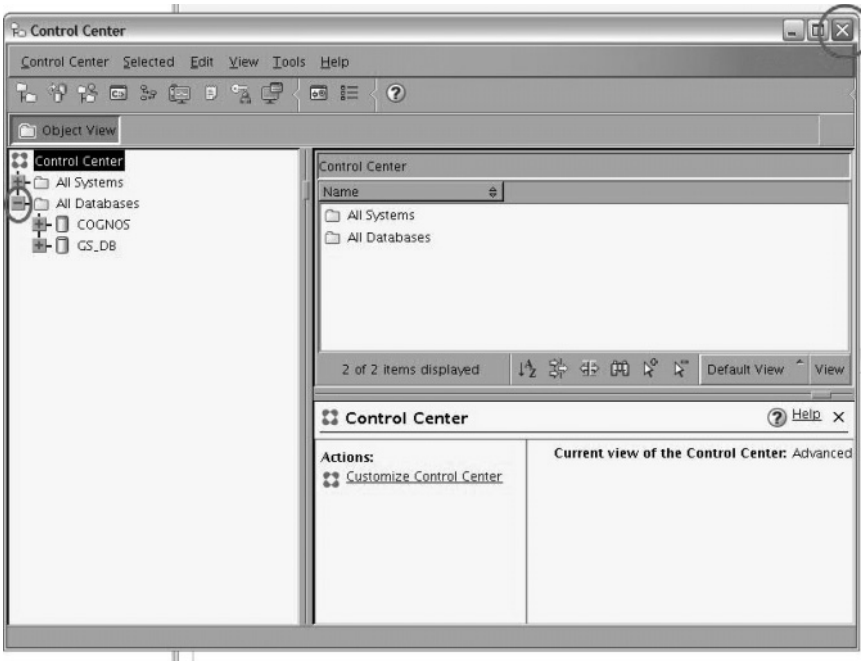

Figure 3.34: Confirming the GS_DB samples database.

11. Restore the contents of the database using the **db2move** utility. Issue the command from within the **GS_DB** folder (or **GS_DB/data**, in Cognos 10), which contains the sample contents:

```
$ cd /opt/ibm/cognos/webcontent/samples/datasources/db2
$ gunzip GS_DB.tar.gz
$ tar -xvf GS_DB.tar
$ cd GS_DB
$ db2move GS_DB import -u db2inst1 -p db2inst1
```

12. Add the appropriate users and schema privileges using the Control Center again:

```
$ db2cc &
```

13. Select “Advanced View” and click **OK**. Under All Databases > GS_DB > User and Group Objects > DB Users, select the “Add New User” link (Figure 3.35).

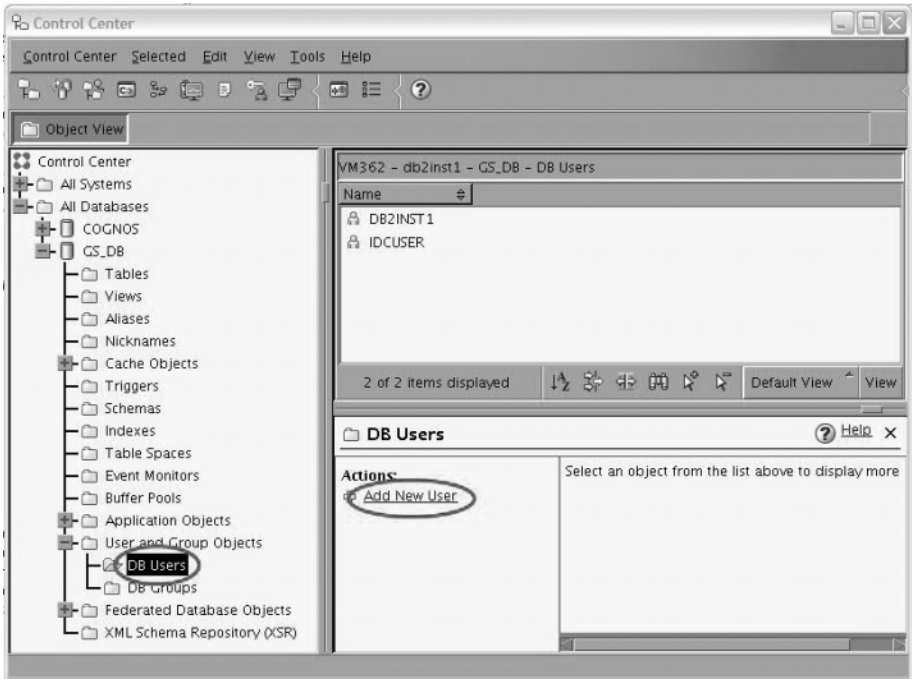

Figure 3.35: Adding a user to the DB2 samples database.

14. Create the **GOSALES** user and allow it to connect to the database, as shown in Figure 3.36. Click **OK**.

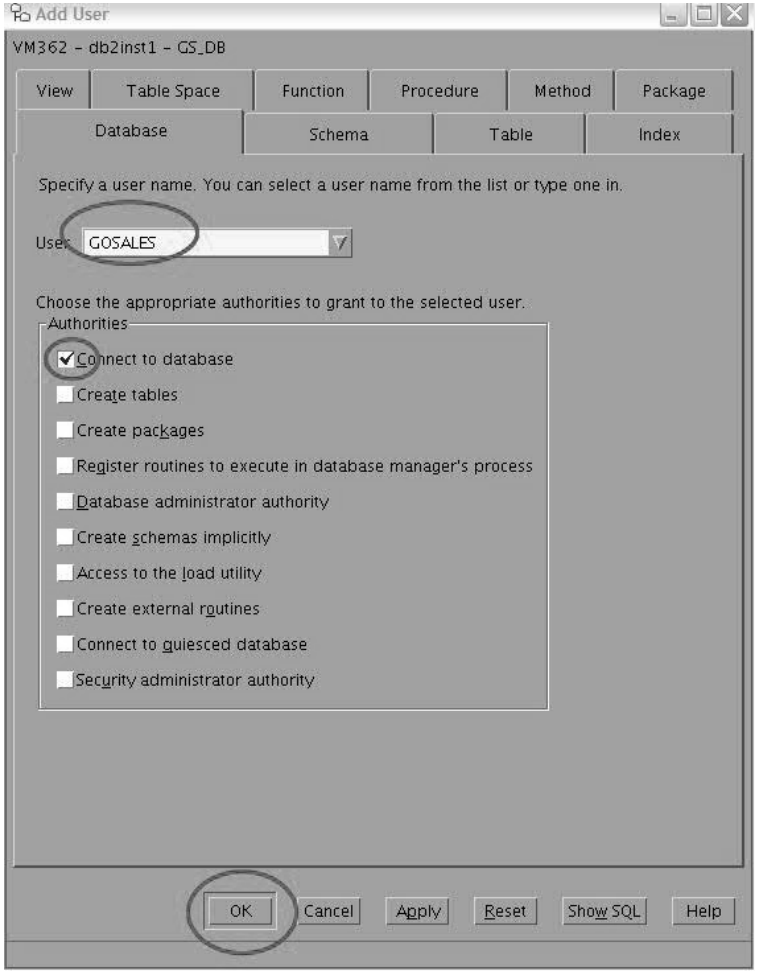

Figure 3.36: Creating the GOSALES user.

15. Create the **GOSALESDW** user and allow it to connect to the database, as shown in Figure 3.37. Click **OK**.

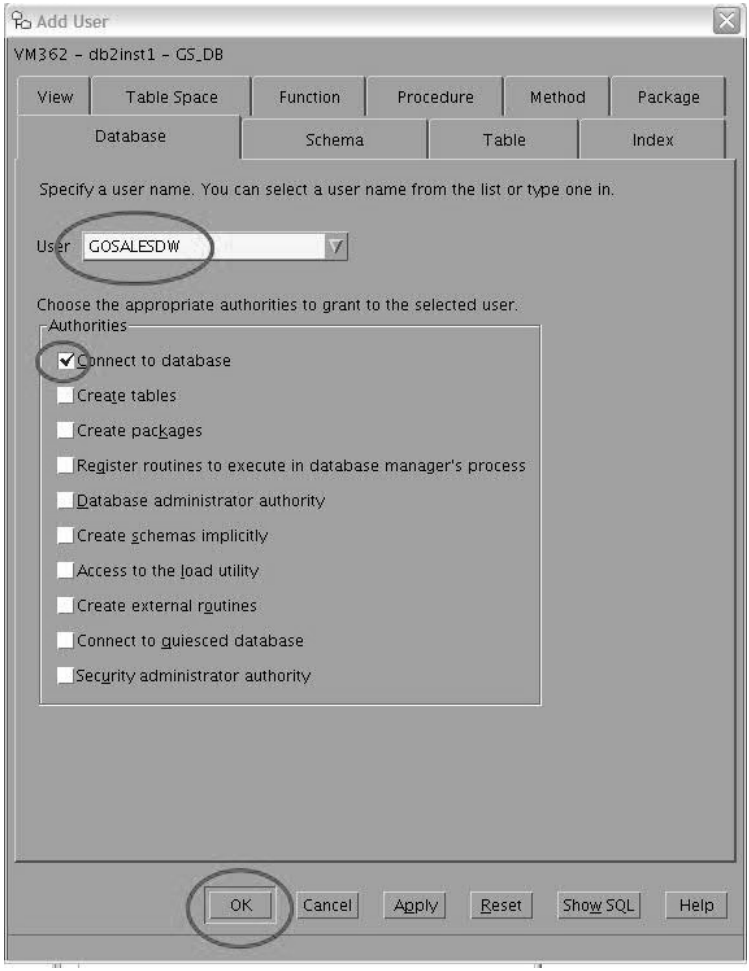

Figure 3.37: Creating the GOSALESDW user.

16. You should now have the **GOSALES** and **GOSALESDW** users created, as shown in Figure 3.38. Close the Control Center.

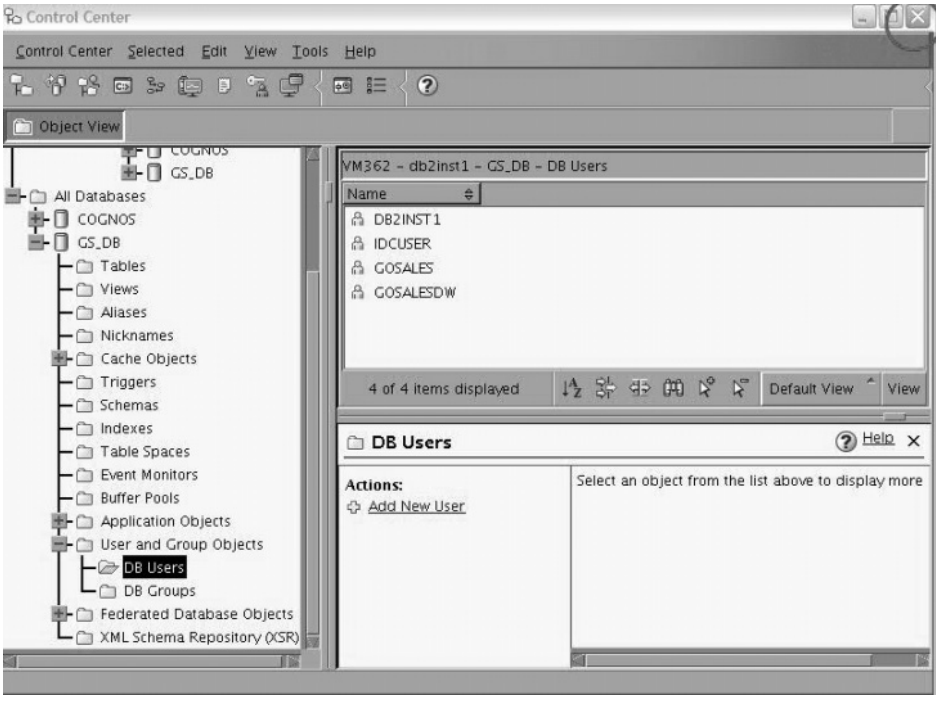

Figure 3.38: Confirming the new DB2 users.

17. Modify the provided **gs_db_modify.sql** script (after making a backup copy) to include the correct DB2 credentials (**db2inst1**):

```
$ cd /opt/ibm/cognos/webcontent/samples/datasources/db2
$ cp gs_db_modify.sql gs_db_modify_orig.sql
$ vi gs_db_modify.sql
```

18. On the first line, replace "<username>" and "<password>" with **db2inst1**.
19. Run the script file, writing the log file to **gs_db_modify.log:**

```
$ rm -f gs_db_modify.log
$ db2 -td@ -f gs_db_modify.sql -z gs_db_modify.log
```

You can verify the contents of **gs_db_modify.log**, if you like.

20. Optionally, confirm the connectivity using the DB2 Configuration Assistant:

```
$ db2ca &
```

Right-click the **GS_DB** database, select "Test Connection" as shown in Figure 3.39, and enter your **db2inst1** credentials.

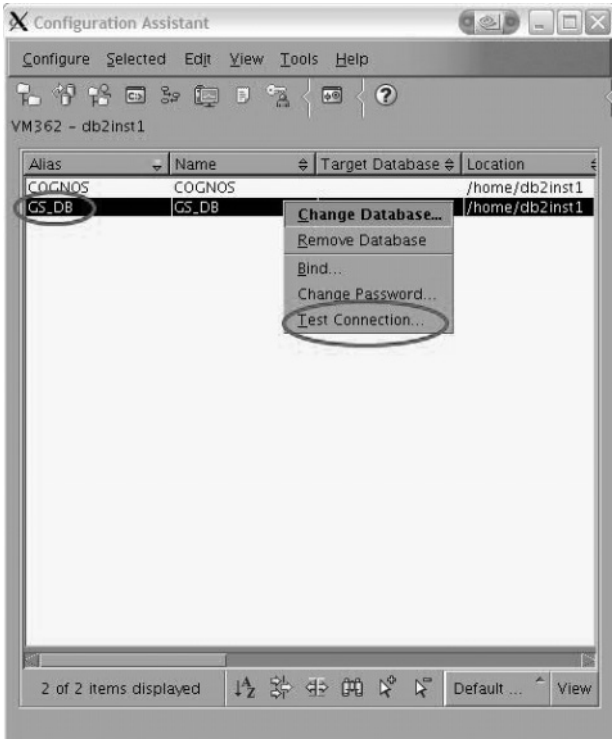

Figure 3.39: Testing the DB2 samples database connection.

Create the Samples Data Source and Deployments in Cognos BI

You are now ready to create the Data Source entry for the DB2 samples in Cognos BI and import the deployment with the sample reports. Do this by following the steps outlined in the Installation Guide included with the Cognos software. In summary, these steps are as follows:

1. Create a new Data Source connection called **great_outdoors_sales** to the DB2 data source **GS_DB**, with a username and password of **db2inst1**.
2. Create a new Data Source Connection called **great_outdoors_warehouse** to the DB2 data source **GS_DB**, with a username and password of **db2inst1**.
3. Copy the provided sample deployment to your Cognos BI installation:

```
$ cp /opt/ibm/cognos/webcontent/samples/content/Cognos_Samples.zip
  /opt/ibm/cognos/deployment
```

4. Using the Content Administration > New Import option, import the samples into your Cognos instance.

Test Your Sample Reports

Congratulations! You now have the sample reports and data deployed to your cloud installation of Cognos BI, as shown in Figure 3.40. You can view, run, or edit the reports.

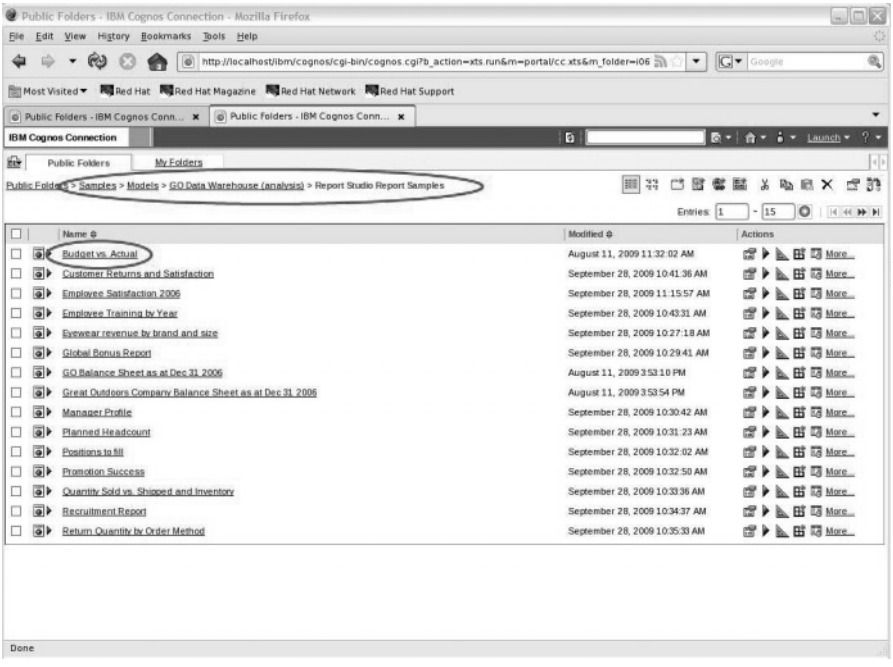

Figure 3.40: A successful deployment of the Cognos BI sample reports.

Configure Security and Access

Now is a good time to do some basic security in the cloud image. There are many best practices around securing your cloud image, including those described in the next section, which merit careful consideration. This is particularly true if there is any sensitivity to the data and reports that will eventually be available on your Cognos installation.

At this point, however, you can take a simple step to limit access to your Cognos installation for computers coming from your company or from specific divisions within your company. Restricting access to your cloud instance will prevent other machines from using the image while it continues to be built and refined. This simple level of security is a good first step as part of an overall security posture and strategy. For some scenarios, it might be sufficient in itself.

Restrict Access to Your Cloud Instance

First, identify all the domains and subdomains that are permitted to access your Cognos instance. For example, you can allow access to computers from within your entire company:

- ibm.com

Or from just certain divisions within your company:

- division1.ibm.com
- division2.ibm.com

Lock Down the Web Server

To lock down the Web server:

1. Load the IBM HTTP Server configuration file into the editor:

```
$ cd /opt/ibm/websphere/httpserver/conf
$ sudo vi httpd.conf
```

2. Add the following lines to the bottom of your file to block all Web access from the public internet except for those from your company, replacing the “allow from” lines with appropriate lines:

```
# Block all web access, except for requests from specific subdivisions
<Location "/">
  order deny,allow
  deny from all
  allow from division1.ibm.com
  allow from division2.ibm.com
</Location>
```

3. Save the file.
4. Restart the server so the new settings are detected:

```
$ sudo /opt/ibm/websphere/httpserver/bin/apachectl restart
```

Now, your Web site will deny access to all browsers except those that come in from machines within approved divisions of your company.

Lock Down the Cognos Platform

You can also ensure that non-Web access to the Cognos platform (such as from SDK applications and other Cognos rich-client tools) is available only from specific domains or subdomains.

1. Start IBM Cognos Configuration on the cloud instance:

```
$ cd /opt/ibm/cognos/bin
$ sudo ./cogconfig.sh
```

2. In the Explorer window, go to Local Configuration > Security > IBM Cognos Application Firewall.
3. Verify that the IBM Cognos Application Firewall is enabled.
4. Click **Edit**. Using the **Add** button, add the host and domains to which access should be restricted in the "Valid domains and hosts" property. Click **OK**.
5. Save the configuration, and restart the services when prompted.

Further Hardening

For most scenarios, the level of simple hardening described in the preceding sections might be sufficient. However, if the data or reports to be transferred to the cloud are highly sensitive, you might need to take some additional steps. Securing a Linux system is a large topic that is beyond the scope of this book, but some possible extra steps to consider at this point include virtual network support and reducing your attack surface.

Whether you choose to continue hardening with these strategies or are comfortable with your current level of security, it is important to consider the overall security strategy. This point is particularly important when graduating from a single image to a more complex, multiple-image topology or from development and test workloads to workloads with more sensitive data. The subsequent section on cloud security provides recommendations and best practices to consider.

Virtual Network Support

Configuring and installing a virtual LAN (VLAN) or Virtual Private Network (VPN) is one of the best ways to secure your instance at this point. These virtual networking technologies permit the cloud instances to securely extend your own corporate network. By default, visibility is restricted to other machines within your network and all traffic flows between secure encryption of traffic.

A virtual network is often best set up by leveraging the built-in facilities on your cloud hosting infrastructure, rather than building it yourself. In most cases, this process involves activating the feature through a Web-based interface and then paying an additional monthly fee. Acquiring this service from your cloud hosting provider guarantees that the VLAN/VPN support is set up correctly for you and automatically patched over time, without any additional maintenance from you. This approach is often the most cost-effective way to secure your cloud instances.

IBM, Amazon, and many of the other major cloud hosts have virtual networking support. Consider this support as part of your cloud security strategy.

Reducing Your Attack Surface

The Cognos cloud instance requires only the small set of ports listed in Table 3.3 to be accessible externally. The remaining ports need not be available. Therefore,

one way to reduce your attack surface is to configure your instance to ignore all network traffic on any other port.

Table 3.3: Externally Accessible Ports	
Port Number	**Function**
80, 443	Browser access to Cognos Connection and other browser applications, such as Report Studio
9080 or 9300	Rich-client and third-party access to the Cognos platform
22	ssh access for command-line, file-transfer, and X-Windows access (PuTTY, WinSCP, and Xming)

In addition, you can configure your instance to ignore all network traffic that comes in from any IP address that does not correspond to a recognized set. Recognized sets are usually expressed as Classless Inter-Domain Routing (CIDR) blocks that are small, relatively fixed, and change very slowly over time.

With both a set of ports and CIDR sets, you can use tools such as **iptables** (built into the Linux operating system) to explicitly drop all network packets that do not correspond to a valid port and IP address. This type of blocking is better than the steps you took previously to lock down the Web server. This is because malicious network traffic that might be probing for target instances and ports will be ignored (as if the target instance doesn't even exist), rather than explicitly denied (which, while secure, acknowledges the existence of the target and therefore might invite further attacks).

If you choose to experiment with address and port blocking at this point as part of your image hardening, be sure to first take a private image snapshot of your image. You might accidentally block access to your image from everyone—including yourself! Having a previous image will let you roll back to a state before those changes, in case this happens.

Connect Windows Rich Clients to Your Cognos Cloud Instance

You can use your cloud instance of Cognos 8 or Cognos 10 as the server for your Windows-based Cognos software clients. This lets you use your cloud Cognos BI Server with programs that can be installed on user machines and not in the IBM Cloud, such as the following:

- Framework Manager
- Go! Office or Cognos for Microsoft Office

- PowerPlay
- Transformer

The process for connecting these Windows programs to the cloud is the same as if the Cognos BI Server were installed on a server within the firewall. In other words, configure the programs with the correct gateway and dispatcher URIs, and ensure that connectivity between the Windows client machines and the server exists.

Here are the main steps involved:

1. Verify the address of your machine instance.
2. Ensure connectivity between the Windows machine and the Cognos dispatcher. The dispatcher is located at your instance's address. By default, this is under port 9080 if you are using WebSphere Application Server or 9300 if you are using Tomcat.
3. Configure the Windows machine with the correct gateway and dispatcher URIs.

To verify the address of your machine instance, open a browser on the Windows machine and point it to the Cognos instance, "http://*vhost.ibm.com*/ibm/cognos" (replacing *vhost.ibm.com* with the address provided to you from the IBM Cloud's Control Panel). You should see the Cognos BI splash screen and main portal page.

A simple way to check for connectivity between the Windows machine and the cloud dispatcher is to use the **telnet** command from the Windows PC. On that machine, run the command prompt, and enter the following command, replacing *vhost.ibm.com* with the server address or IP address of the cloud instance and 9080 with the correct dispatcher port, if necessary:

```
> telnet vhost.ibm.com 9080
```

The screen should clear, and you should be able to type directly to the Cognos dispatcher. If you see this, you have successfully connected from the Windows machine to your cloud's Cognos BI dispatcher. You will not see any characters or be able to have a meaningful authentication session with the dispatcher, so at this point simply press **CTRL+]** and type **quit** to exit the telnet application.

If you receive a message that says, "Could not open connection to the host, on port 9080: Connect failed," there is no connection available. You will need to diagnose the problem further. Common issues to verify include the following:

- Ensure that the Cognos BI instance is running and that the address and port values are correct.
- Ensure that the port has public access enabled. Refer to the documentation in the IBM Smart Business Developer and Test Cloud to see which ports are open by default.
- Ensure that you do not have a security configuration on the cloud instance that is blocking this port. For example, check the **iptables** configuration.

Once you have verified connectivity, configuring the Windows software with the correct URIs is the next and final step.

Install the desired Cognos software, such as Framework Manager, using the Windows-based Cognos installer as usual, but omitting any server components in the installation. For Framework Manager, PowerPlay, and Transformer, take these steps:

1. Run Start > IBM Cognos 10 (or 8) > IBM Cognos Configuration to launch the IBM Cognos Configuration tool.
2. Go to Local Configuration > Environment.
3. Replace all instances of *localhost*, as shown in Figure 3.41, with the IP address of the cloud instance. Also replace all instances of 9300 with the dispatcher port of the cloud instance.
4. Select File > Save.

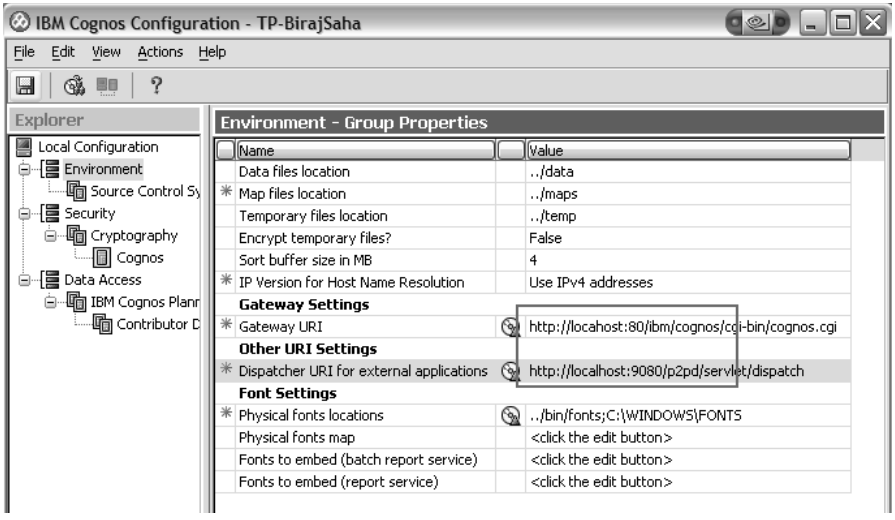

Figure 3.41: Configuring Windows applications to use the Cognos BI Server cloud instance.

For Go! Office or Cognos for Microsoft Office, perform these steps:

1. After installing the software, start Microsoft Office and locate the IBM Cognos toolbar.
2. Click the **Options** button, and locate the IBM Cognos tab in the left navigation pane.
3. Under IBM Cognos Systems, in the "System gateway" URI box, enter the gateway address of your cloud instance. For example, enter "http://*vhost.ibm.com*/ibm/cognos/cgi-bin/cognos.cgi," replacing *vhost.ibm.com* with the cloud instance's address. You will be able to provide a friendly name for the gateway address, and use the **Test Connection** button to ensure that the server is accessible.

Create a Cognos BI Cloud Image

Once you have the software configured to your liking, it is a good time to create a private image of your Cognos 8 or Cognos 10 BI Server cloud instance. An image represents a snapshot of the instance that can be downloaded and reused to create new instances. The objective is to ensure that the image you create can generate instances of fully configured Cognos BI installations, ready for immediate use.

To convert your instance into a reusable image, you need to address two items:

- The Cognos content store database is directly dependent on the IP address of the current host and thus cannot be merely cloned and reused on another host. Cognos' deployment archives will be used to reimport and rebuild the content store.
- The cloud instance's host name is hardcoded in various areas of each product, so you will need to update your **hosts** file to treat that host name as "localhost."

Follow the additional setup operations listed below to create the Cognos cloud instance. Once an image is created for it, any new instance based on the image will have Cognos BI Server running automatically, available immediately after the new instance has been created!

Export Your Cognos Sample Public Folders and Directory Content

In Cognos Administration, create a new deployment export containing all public folders and directory content. Include all packages, data source connections, and sign-ons, with an encryption password. This will create a deployment archive whose content can be readily imported when Cognos BI and WebSphere are first started on new cloud instances created from the image.

Follow these steps:

1. In the "Content_Administration" section, click the **New Export** toolbar button, shown in Figure 3.42.

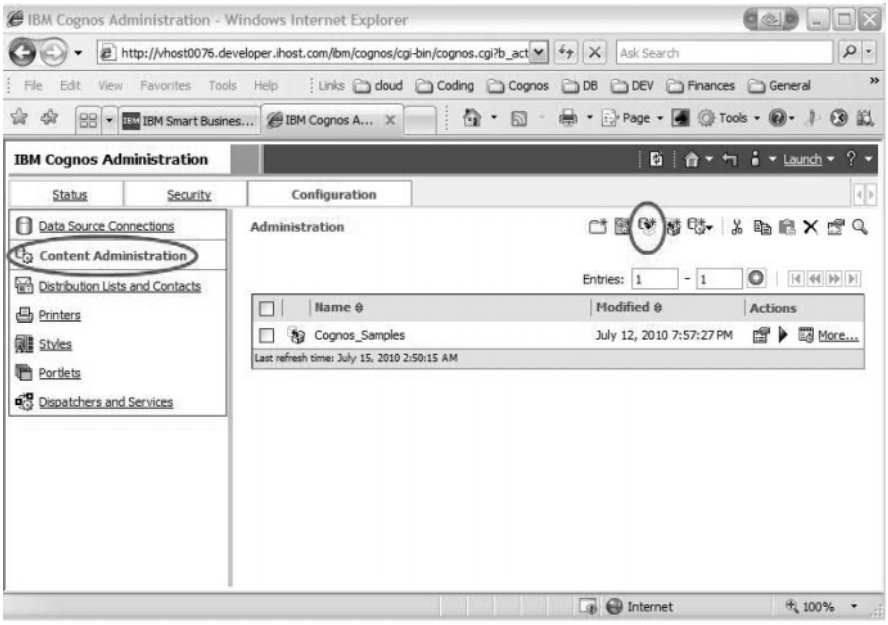

Figure 3.42: Exporting your Cognos content.

2. In the "Specify a name and description" page (Figure 3.43), enter the name **Cognos_Export**, and click **Next**.

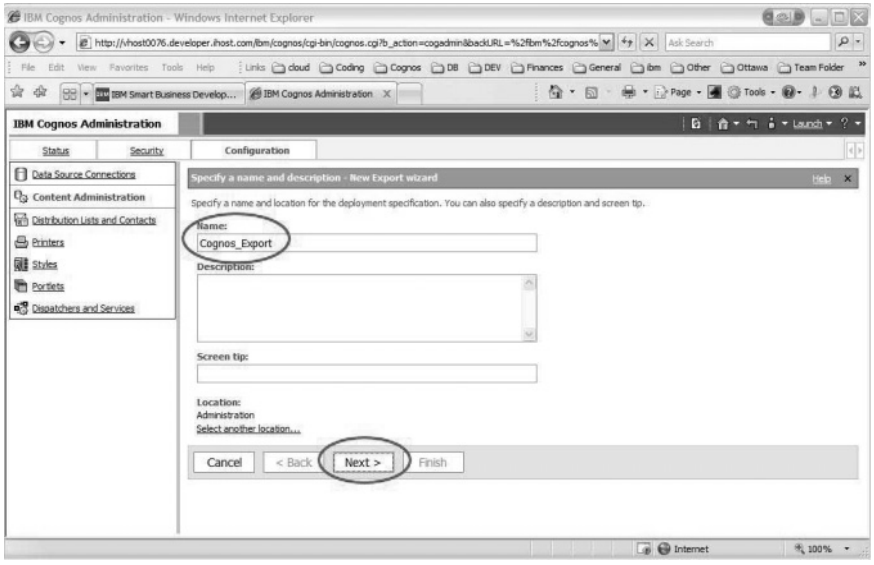

Figure 3.43: Naming the Cognos export job.

3. In most cases, the default option "Select public folders and directory content" is correct. If that is true for you, click **Next**, as shown in Figure 3.44. If you have added additional content to the instance, such as users, and want to include that content in your Cognos BI image, use the "Select the entire Content Store" option.

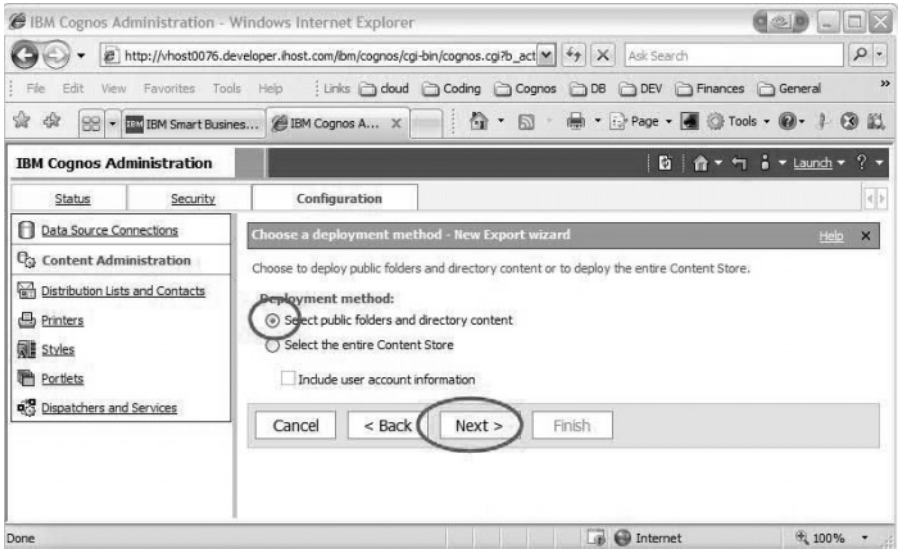

Figure 3.44: Specifying the type of content to export.

4. Select the content to include, and click **Next**. The example in Figure 3.45 shows the "Samples" folder selected.

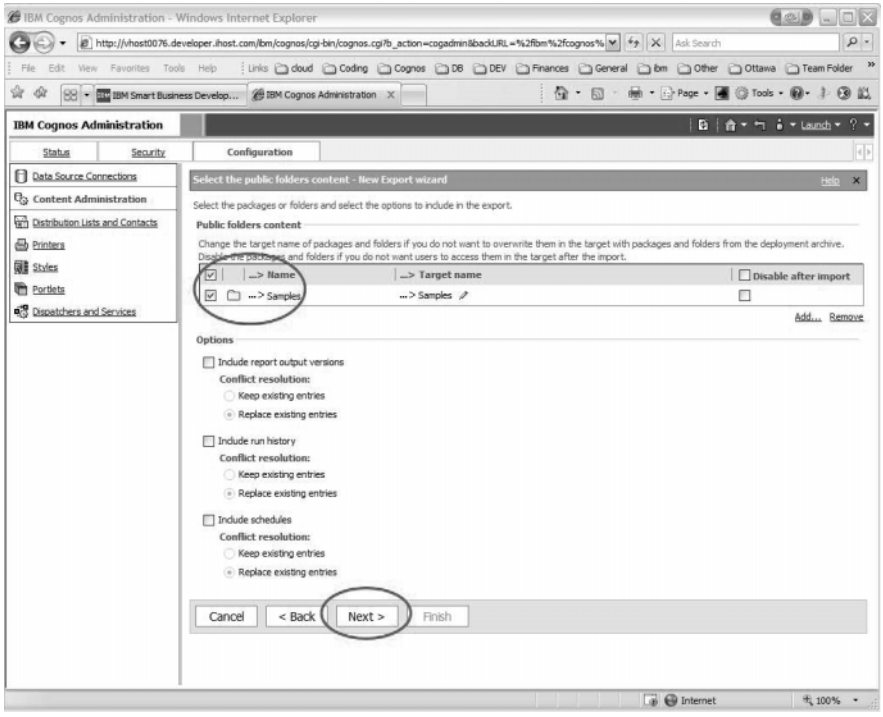

Figure 3.45: Selecting the content to export.

5. Check the “Include data sources and connections” and “Include signons” boxes (Figure 3.46). Click **Next**.

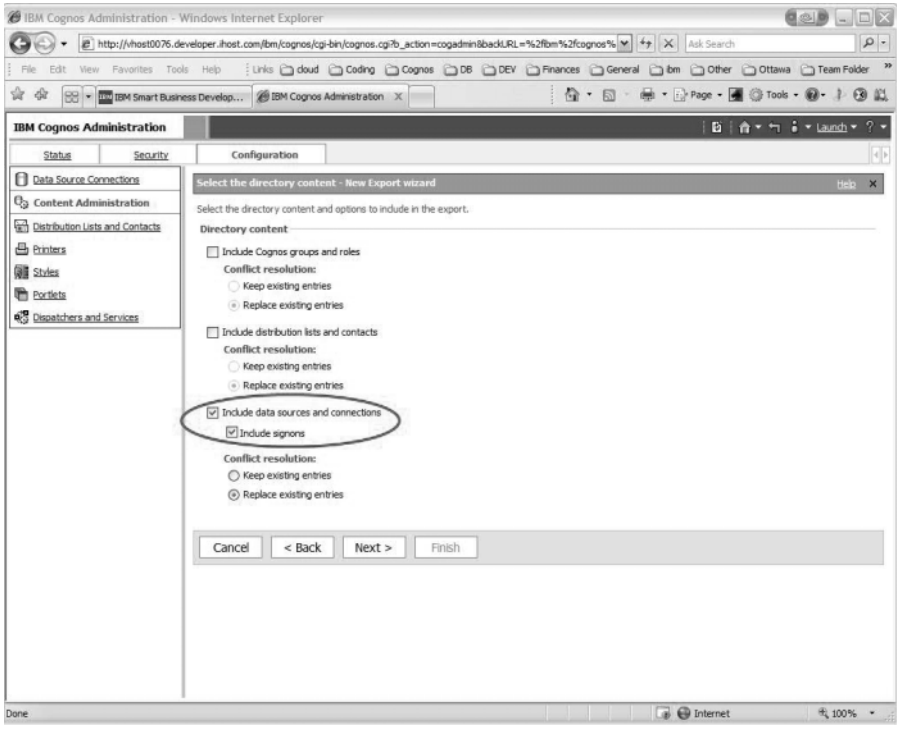

Figure 3.46: Including sign-on information with the export.

6. On the "Specify the general options" page (Figure 3.47), accept the defaults and click **Next**.

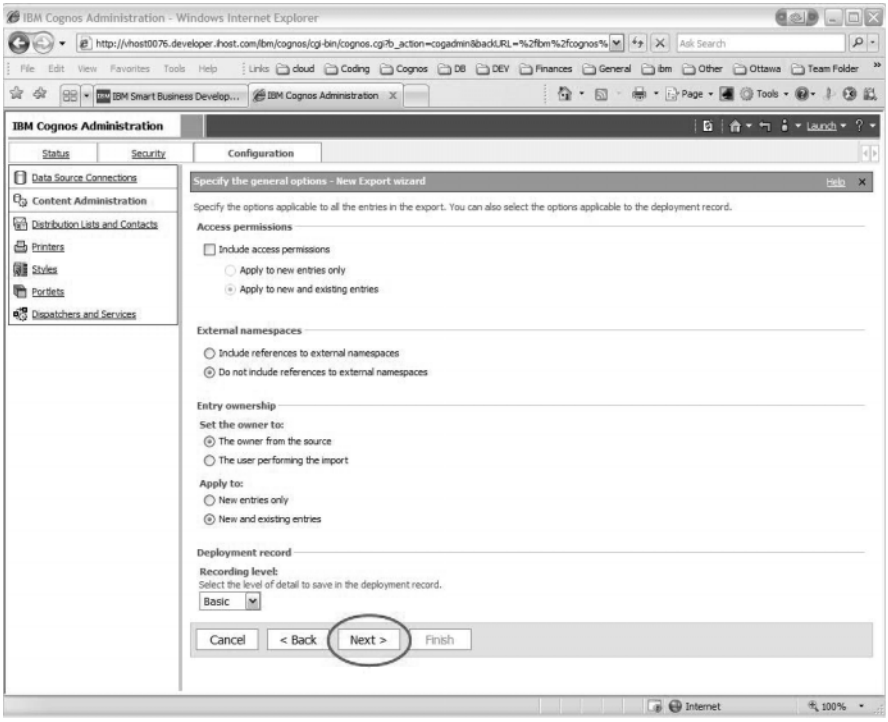

Figure 3.47: Verifying other export options.

7. In the "Specify a deployment archive" page (Figure 3.48), select "Edit the encryption password," set the password to "Cognos_Export," and click **Next**.

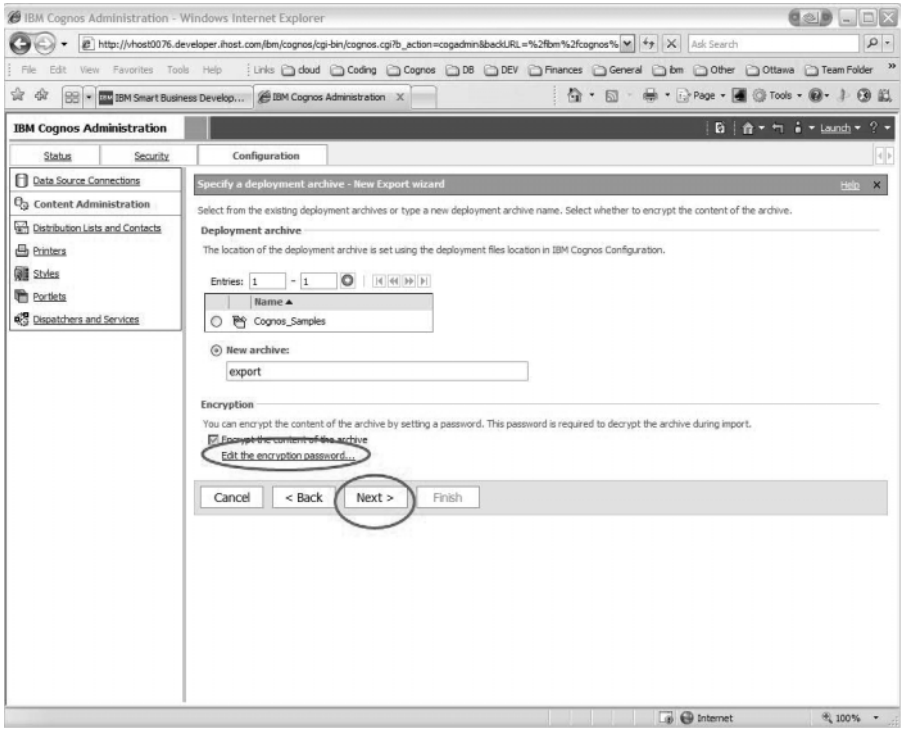

Figure 3.48: Naming the export archive and setting an encryption password.

8. Review the settings on the "Review the summary" page, and click **Next**.

9. On the “Select an action” page (Figure 3.49), select “Save and run once” and click **Finish**.

Figure 3.49: Saving and running the export job.

10. On the "Run with options" page (Figure 3.50), click **Run**. Then click **OK** to perform the export.

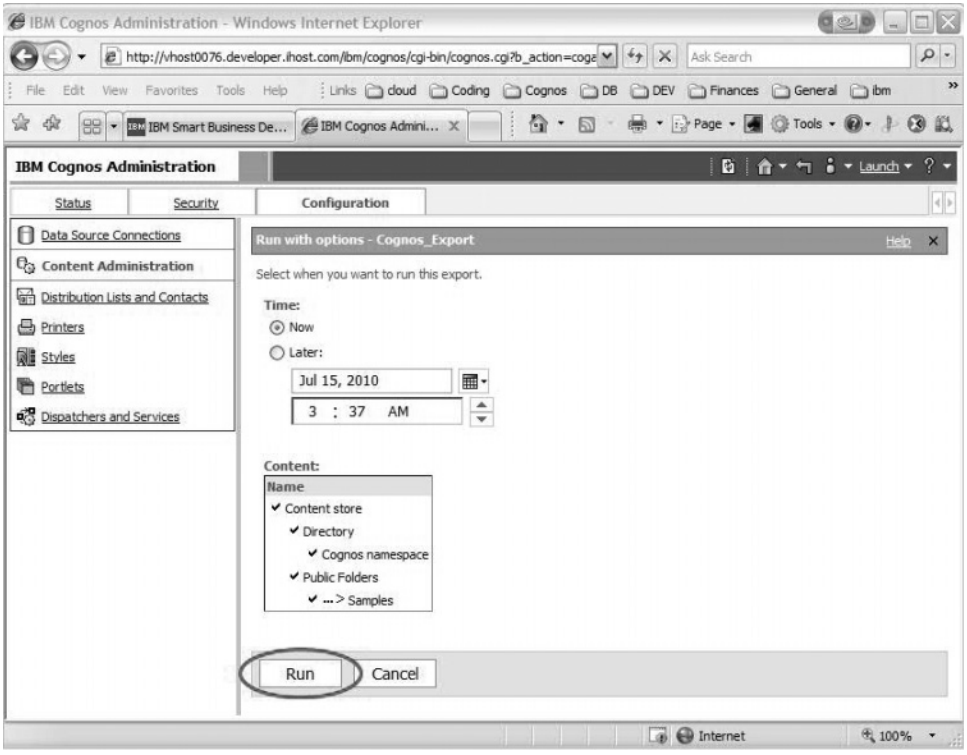

Figure 3.50: Running the export job.

11. Verify the successful export by navigating to the "Past Activities" section and looking for a "Succeeded" status (Figure 3.51).

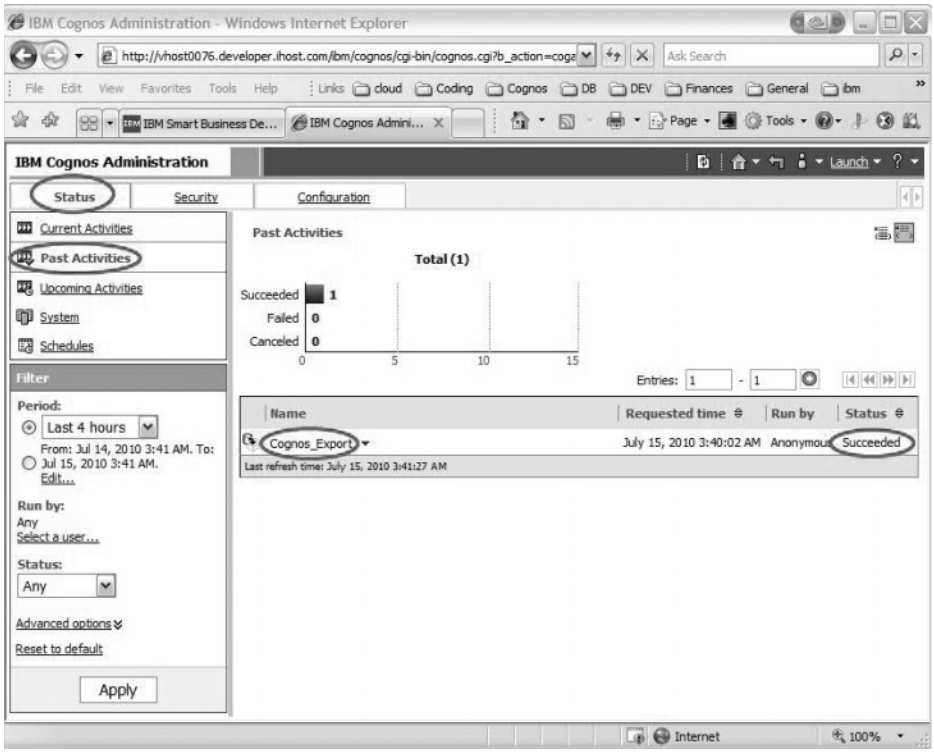

Figure 3.51: A successful export of the Cognos content.

Export Your Cognos Configuration Settings

Cognos Configuration can export the instance's configuration to a file called **export.xml**. This file will be used to automate the configuration of new instances of Cognos BI upon creation. Note that this file stores the username and password for the DB2 content store unencrypted.

1. Run Cognos Configuration, and select File > Export As, as shown in Figure 3.52. In the warning dialog that appears, click **Yes**.

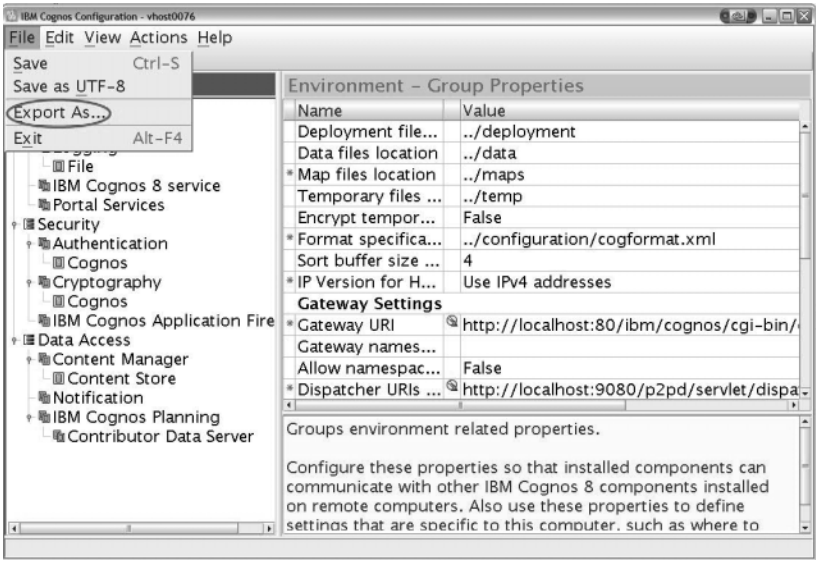

Figure 3.52: Exporting your Cognos configuration.

2. In the Export As dialog, enter **export** as the file name, and click **Save** (Figure 3.53).

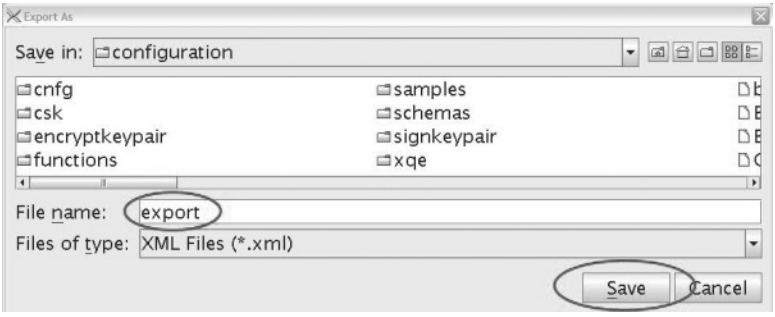

Figure 3.53: Saving the exported configuration as an XML file.

3. After the task execution dialog confirms that all tasks have been completed successfully, click **Close** and exit the program.
4. At this point, an **export.xml** file, located in the **configuration** directory, has been created and contains details of the instance's Cognos configuration. Make a copy of this directory to one called **configuration_master**. We'll use the file in the **cognos_init.sh** script we create in the next section.

```
$ cp -rf /opt/ibm/cognos/configuration
  /opt/ibm/cognos/configuration_master
```

Prepare Image Configuration Scripts

When a new instance is created from an image, the IBM Cloud allows Unix scripts to be automatically executed upon startup. Use this feature, along with configuration scripts, to prepare the Cognos software for use whenever your image is used to create new instances. These scripts will do the following:

- Automatically re-create the content store database upon instance startup
- Launch Cognos BI Server and its underlying services: WebSphere Application Server, IBM HTTP Server, and DB2 Server

The administrator can use these scripts after the new instance is created to automatically stop and start the installed services.

1. Build a script called **cognos_stop.sh** to stop all services:

```
$ vi /home/idcuser/cognos_stop.sh
```

   This script can be used by administrators of the new instance. It will also be invoked by the **/etc/init.d/cloud-startup3.sh** script below when newly created instances based on your image are stopped.
2. Add the following lines to the script:

```
#
# cognos_stop.sh
#  - intended to be invoked whenever Cognos BI needs to be
#    stopped, e.g., when the cloud instance is “restarted”.
#
DB2DIR=/opt/ibm/db2; export DB2DIR
DB2INSTANCE=db2inst1; export DB2INSTANCE
PATH=$PATH:$HOME/bin:$DB2DIR/bin; export PATH
LD_LIBRARY_PATH=/opt/ibm/cognos/bin:$DB2DIR/lib32;
  export LD_LIBRARY_PATH
sudo ksh $DB2DIR/cfg/db2profile
JAVA_HOME=/opt/ibm/websphere/appserver/java; export JAVA_HOME
PATH=$JAVA_HOME/jre/bin:$PATH; export PATH
LD_PRELOAD=/usr/lib/libfreetype.so.6; export LD_PRELOAD
DB2CODEPAGE=1208; export DB2CODEPAGE
sudo /opt/ibm/websphere/httpserver/bin/apachectl stop
sudo chmod -Rf g+wx /opt/ibm/websphere/appserver
```

```
/opt/ibm/websphere/appserver/bin/stopServer.sh server1
  -username idcuser -password idcuser
db2 force application all
db2 db2stop
```

3. Ensure that the script is executable:

```
$ chmod +x /home/idcuser/cognos_stop.sh
```

4. Build a script called **cognos_start.sh** to start all services. This script will also be used by the **/etc/init.d/cloud-startup3.sh** script below when newly created instances based on your image are started.

```
$ vi /home/idcuser/cognos_start.sh
```

This script can be used by administrators of the new instance.

5. Add the following lines to the script:

```
# cognos_start.sh
#  - intended to be invoked whenever Cognos BI needs to be
#    started, e.g., when the cloud instance is “restarted”.
#
DB2DIR=/opt/ibm/db2; export DB2DIR
DB2INSTANCE=db2inst1; export DB2INSTANCE
PATH=$PATH:$HOME/bin:$DB2DIR/bin; export PATH
# echo NOTE!! For 64-bit installs, only change the value of
  “/opt/ibm/cognos/bin” to “/opt/ibm/cognos/bin64”. However,
  do not change $DB2DIR/lib32.
LD_LIBRARY_PATH=/opt/ibm/cognos/bin:$DB2DIR/lib32;
  export LD_LIBRARY_PATH
sudo ksh $DB2DIR/cfg/db2profile
JAVA_HOME=/opt/ibm/websphere/appserver/java; export JAVA_HOME
PATH=$JAVA_HOME/jre/bin:$PATH; export PATH
LD_PRELOAD=/usr/lib/libfreetype.so.6; export LD_PRELOAD
DB2CODEPAGE=1208; export DB2CODEPAGE
db2 db2start
sudo /opt/ibm/websphere/httpserver/bin/apachectl start
sudo chmod -Rf g+wx /opt/ibm/websphere/appserver
/opt/ibm/websphere/appserver/bin/startServer.sh server1
  -username idcuser -password idcuser
```

For a 64-bit install, use the following line in the script instead:

```
LD_LIBRARY_PATH=/opt/ibm/cognos/bin64:/$DB2DIR/lib64
```

6. Ensure that the script is executable:

```
$ chmod +x /home/idcuser/cognos_start.sh
```

7. Create a script called **cognos_init.sh** that will also be invoked by **/etc/init.d/cloud-startup3.sh** when your image is used to create a new instance, upon first startup. This script re-creates the underlying Content Store database and performs a silent reconfiguration of the Cognos BI Server.

```
$ sudo vi /home/idcuser/cognos_init.sh
```

Add the following lines to the script:

```
#
# cognos_init.sh
#  - intended to be invoked by root user
#  only one time on initial instance startup
#
echo Setting up required environment variables
PATH=$PATH:$HOME/bin; export PATH
DB2DIR=/opt/ibm/db2; export DB2DIR
DB2INSTANCE=db2inst1; export DB2INSTANCE
PATH=$PATH:$DB2DIR/bin; export PATH
LD_PRELOAD=/usr/lib/libfreetype.so.6; export LD_PRELOAD
LD_LIBRARY_PATH=/opt/ibm/cognos/bin:$DB2DIR/lib32;
  export LD_LIBRARY_PATH
sudo ksh $DB2DIR/cfg/db2profile
JAVA_HOME=/opt/ibm/websphere/appserver/java; export JAVA_HOME
PATH=$JAVA_HOME/jre/bin:$PATH; export PATH
DB2CODEPAGE=1208; export DB2CODEPAGE

echo “Stopping the services just in case any are still running”
/home/idcuser/cognos_stop.sh

echo “Running the script to recreate the content store database
 via DB2 CLP”
db2 -td@ -f /home/idcuser/cognos_init_recreateContentStore.sql
```

```
echo “Clean existing logs”
rm -rf /opt/ibm/cognos/logs/*
  /opt/ibm/websphere/appserver/profiles/AppSrv01/logs/server1/*

echo “Reconfigure Cognos BI Server via silent config”
rm -rf /opt/ibm/cognos/configuration
cp -rf /opt/ibm/cognos/configuration_bak
  /opt/ibm/cognos/configuration
cp /opt/ibm/cognos/configuration_master/export.xml
  /opt/ibm/cognos/configuration/cogstartup.xml

echo “--- starting up DB2”
db2 db2start
echo “--- running Cognos BI silent config”
  /opt/ibm/cognos/bin/cogconfig.sh -s
echo Starting DB2 server
db2 db2start

echo Starting IBM HTTP server
  /opt/ibm/websphere/httpserver/bin/apachectl start

echo Clean existing logs
rm -rf /opt/ibm/cognos/logs/*
  /opt/ibm/websphere/appserver/profiles/AppSrv01/logs/server1/*

echo Reconfigure Cognos via silent config
cp -rf /opt/ibm/cognos/configuration
  /opt/ibm/cognos/configuration_master
cp /opt/ibm/cognos/configuration_master/export.xml
   /opt/ibm/cognos/configuration/cogstartup.xml
   /opt/ibm/cognos/bin/cogconfig.sh -s

echo Starting Cognos via WebSphere Application Server
/opt/ibm/websphere/appserver/bin/startServer.sh server1
  -username idcuser -password idcuser
```

For a 64-bit install, use the following lines instead:

```
LD_LIBRARY_PATH=$DB2DIR/lib64; export LD_LIBRARY_PATH
LD_PRELOAD=/usr/lib64/libfreetype.so.6; export LD_PRELOAD
LD_LIBRARY_PATH=/opt/ibm/cognos/bin64:/$DB2DIR/lib64
```

8. Ensure that the script is executable:

```
$ sudo chmod +x cognos_init.sh
```

9. As the root user, edit the cloud startup script **/etc/init.d/cloud_startup3.sh.** This script will stop any currently running services, prepare the new instance, and then invoke the scripts built above to rebuild, reconfigure, and then start up the IBM Cognos BI Server.

```
$ sudo vi /etc/init.d/cloud-startup3.sh
```

Replace the file with the contents below. Ensure you replace the occurrence of **vhost** below with your cloud master instance's hostname.

```
### BEGIN INIT INFO
# chkconfig 3 90 20
# description: cloud-startup3.sh
# processname: cloud-startup3.sh
# Provides:         cloud-startup3.sh
# Required-Start:
# Should-Start:
# Required-Stop:
# Should-Stop:
# Default-Start:    3
# Default-Stop:
# Short-Description: Cloud startup
# Description:      Extract and set user password
### END INIT INFO

case "$1" in
  start)
    HOSTNAME_MASTER=vhost
    echo "== Cloud Starting, master instance: $HOSTNAME_MASTER"
    if [ ! -e /etc/cloud/idcuser_pw_randomized ]; then
       echo "Randomizing idcuser password"
       echo idcuser:'< /dev/urandom tr -dc _A-Z-a-z-0-9 |
         head -c16' | /usr/sbin/chpasswd
       touch /etc/cloud/idcuser_pw_randomized
    fi
    echo "== Preparing Cognos BI Server"
```

```
    if [ 'hostname' == $HOSTNAME_MASTER ]
    then
      echo "== Master Instance"
      sudo -u idcuser /home/idcuser/cognos_stop.sh
    else
      echo "== Image Instance"
      if [ -s "/home/idcuser/cognos_init.sh" ] ; then
           echo "== Initializing Cognos (first time only)"
           echo "Initial cleanup"
           history -c
           sudo chmod g+w /var/log/lastlog
           echo > /var/log/lastlog
           chmod g-w /var/log/lastlog
           echo "Adjusting the hostname in DB2"
           chmod g+w /home/db2inst1/sqllib/db2nodes.cfg
           echo 0 $HOSTNAME 0 >
             /home/db2inst1/sqllib/db2nodes.cfg
           chmod g-w /home/db2inst1/sqllib/db2nodes.cfg
           echo "Adjusting the /etc/hosts file to map image
             master instance hostname to localhost"
           chmod g+w /etc/hosts
           echo 127.0.0.1 $HOSTNAME_MASTER >> /etc/hosts
           chmod g-w /etc/hosts
           echo "running cognos_init script to recreate the
             content store"
           sudo -u idcuser /home/idcuser/cognos_init.sh >
             /home/idcuser/.cognos_init.log
           sudo -u idcuser mv /home/idcuser/cognos_init.sh
            /home/idcuser/.cognos_init.sh
      fi
    fi
    echo "== Starting Cognos BI Server"
    sudo -u idcuser /home/idcuser/cognos_start.sh > /home/idcuser/.
   cognos_start.log
    ;;
  stop)
    echo "== Cloud Stopping"
    sudo -u idcuser /home/idcuser/cognos_stop.sh
    ;;
  *)
    echo "Usage: $0 {start|stop}"
    exit 1
    ;;
esac
```

10. At this point, before creating your image, you should remove sensitive or extraneous files and data. Here are some recommended commands to run:

```
# Clean logs
$ rm –rf /opt/ibm/cognos/logs/*
  /opt/ibm/websphere/appserver/profiles/AppSvor01/logs/server1/*

# Erase “last login” information
$ sudo chmod g+w /var/log/lastlog
$ sudo echo > /var/log/lastlog
$ sudo chmod g-w /var/log/lastlog

# Remove the installations and tarballs
$ rm –rf /home/idcuser/installs

# Clear command history
$ history -c
```

Create the IBM Cloud Private Image

You can now create the image on the IBM Smart Business Development and Test for the IBM Cloud from your IBM Cloud Control Panel by using your instance’s Create Image link, as shown in Figure 3.54.

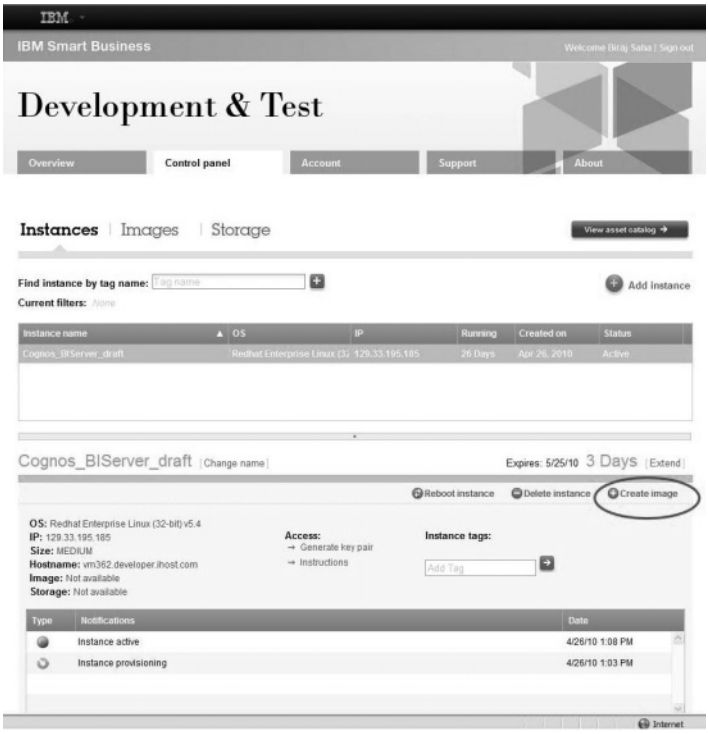

Figure 3.54: Creating a new private image.

Follow the provided instructions to complete the image creation. The image creation process can take several hours to complete. When it's completed, the new image will appear with a status of "Available" in your Control Panel's "Images" section.

Congratulations! You now have a private image that you can use to create and deploy new Cognos instances in minutes, rather than days!

Installation Variations

Cognos 8 and Cognos 10 Business Intelligence Server provide a number of different installation options. For readability, this chapter provided a step-by-step guide based on a standard installation path. However, you might want to explore other variations in your cloud image:

- **64-bit installation**. The steps provided in this chapter are focused on a 32-bit installation, but variations are provided for a 64-bit installation.

Large installations of Cognos BI can benefit from the larger memory space available in 64-bit environments.

- **Non-root installation**. WebSphere and DB2 were both installed using a traditional method—as a user with root privileges. However, installing as a non-root user has the benefit of better containing the application and ensuring that root privileges are not needed to run the application. At present, though, non-root installations of DB2 result in limitations that might be undesirable, such as only one instance per user and some missing tools (e.g., **db2ca**, **dbrcc**).
- **Application and Web server variations**. This chapter used WebSphere Application Server as the application server and IBM HTTP Server as the Web server. However, many other choices exist. Tomcat Web Application Server is supplied as part of the Cognos BI installation, and Apache HTTP Server is often used as well. Another variation is to remove the Web server entirely and have clients connect directly to WebSphere Application Server (i.e., have WebSphere act as both the application server and the Web server).
- **Alternative desktop clients**. Rather than using Xming as a local X-Windows client, you can leverage other tools, such as Cygwin/X, Exceed, VNC, and XWin32.
- **Alternative content stores**. Rather than DB2, other databases, or Cognos's included Cognos Content Store, can be used.

In all cases, as you plan your installation, consider the requirements of your workload to evaluate these and other variations. For example, using Tomcat and the Cognos Content Store might provide a much simpler installation and smaller image for your small developer and test workload. However, it might not be appropriate for a heavier or more demanding workload.

4

Security Best Practices

On-premises security is still applicable to the cloud. Cognos BI provides a security architecture that enables you to secure the application, including features such as the IBM Cognos Application Firewall, the cryptographic environment, security temporary files, and access to Cognos Connection and data sources. On the cloud, these features continue to be an important part of your security plan. For descriptions of these features and associated best practices, see the "Security" chapter of the *IBM Cognos Business Intelligence Architecture and Deployment Guide* (available at *http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.crn_arch.8.4.0.doc/crn_arch.html*).

In the cloud, it becomes critically important to secure the operating system and the deployment. At least in theory, cloud instances in your solution are exposed to the public internet and therefore might be at greater risk for attacks than an internal installation. This exposure and risk apply not just to the cloud instance, but to any data in transit between your on-premises environment and the cloud.

The risk of attack, confidentiality loss, data integrity loss, and data theft is not unique to cloud deployments. These risks have always existed, even within the most securely firewalled data center (e.g., through employee theft). However,

the public use of the cloud, along with its complex networking and shared environment, makes security deployment especially critical.

The *IBM Cognos Business Intelligence Architecture and Deployment Guide* makes several security recommendations, including best practices regarding the following:

- User accounts
- Account policies
- File system permissions
- Network services
- System patches
- Operating system minimization
- Logging and monitoring
- System integrity
- Firewall and network ports
- Web server configuration

All the recommendations provided in the guide remain applicable and relevant in the cloud.

Cloud Security Best Practices

Security practices specific to cloud environments constitute a big topic and an ongoing process that is continuously evolving. Most important, you should assume that until machine instances on the public cloud are secured, they are wide open and subject to attack by both insiders (people to whom you have given explicit access to the cloud instance) *and* outsiders (people from the outside world).

Securing your cloud instances and applications is a mandatory step to protect your users and your investment. The rest of this chapter provides ideas to help you in this process.

Securing the Instance

Install a firewall on every machine instance in your cloud, and lock it down to allow only essential ports to go through. Port filtering is common in a cloud environment.

Here are some additional specific ideas for securing the instance:

- Prevent access to any open ports, except from instances from your own cloud environment. Limiting access to a set of IP addresses is even better.
- Shut down all non-essential services.
- Audit all login attempts, and generate alerts to the administrator if abnormal behavior is detected.

There is also a more sophisticated and advanced recommendation you might want to consider: run all your application in a **chroot** environment. This option requires more effort but results in a much more secure environment.

Securing the Application

Here are some specific ideas for securing the application:

- Revoke all access by default, and grant access only as needed.
- Install a firewall, and limit application-port access to a group of known machines on your cloud environment.
- Limit application access to a set of elastic IPs, as you've secured the virtual machine.
- Audit all application login attempts, and alarm for abnormal behavior.
- Audit all major application events.

Securing the Database

Here are some specific ideas for securing the database:

- Revoke all access by default, and grant access only as needed.
- Prevent data access from the system DBA. Doing so will prevent anyone with DBA privileges (whether legitimate or not) from being able to read the actual data.
- Limit access to database server ports to a group of known machines on your cloud environment.
- Limit database access to a set of elastic IPs, similar to the way you've secured the virtual machine.
- Audit all major database events.

5

Handling Cloud Topologies

We've offered guidance for setting up your Cognos solution into a single image on the cloud. However, just as a traditional data center deployment of Cognos may require multiple machines, your cloud solution may require multiple images. Criteria such as performance, scalability, and high availability typically lead to a multi-image topology. Best practices for these multi-image topologies are described in the following section.

Using the Hosts File to Manage Multiple Images

When dealing with multiple cloud images, you also must deal with multiple and changing IP addresses associated with those images. A hostname pairs a logical name with an IP address. It is often stored by a DNS server, letting programs map the hostname to an IP address, or vice versa. For example, you might map the IP address for your standby Content Manager, 10.3.0.1, to the hostname **cm_standby**.

For complex solutions, it could be appropriate to deploy an internal DNS server to manage these hostnames. However, machines also have a file known as the **hosts** file that the operating system uses to map hostnames to IP addresses. In many cases, the **hosts** file is a simpler way to manage your cluster of instances.

On Unix, the **hosts** file is located at **/etc/hosts**. On Windows, it is located at **\Windows\System32\Drivers\etc\hosts**.

We will assume that all your instances are in the same cloud region and therefore are in the same private class IP range. If you have cloud instances in different geographical regions, you might have the added consideration of dealing with multiple class IP ranges.

Example: An Elastic Cognos Cluster with a Single Image

It is simpler to manage a cloud solution that has as few images as possible. For example, although a Cognos BI cloud deployment may have instances with different roles (e.g., gateway, dispatcher, content manager), it is often easier to manage the deployment by creating a single generic image, which can then be customized.

A Cognos BI cluster may include three logical components: a gateway, a dispatcher and content manager, and a database. One approach to building this cluster would be to use three separate cloud images: a gateway image (with Cognos and Apache), a dispatcher cloud image, and a DB2 database image. However, as mentioned, it is simpler to manage a cloud solution that has as few images as possible. Scaling the solution dynamically in this scenario also presents challenges that are absent when all the components are loaded on a single cloud image.

You can start with a single cloud image that includes all the software components (Cognos BI, Apache, and DB2). You can then define a mechanism to dynamically start instances from this single image and configure it to dynamically grow.

The main steps are as follows:

1. Install all the required software into a single cloud image.
2. Configure the solution cluster, using hostname aliases.
3. Map these aliases into an actual instance and its associated IP address, using the **hosts** file.

Hostname Aliases

We create a **hosts** file that contains the entries listed in Table 5.1.

Table 5.1: Entries in the Hosts File

Hostname alias	Description
myCognos	Alias to the current instance (similar to *localhost*)
vm-db2	Alias to the DB2 instance
vm-gateway	Alias to the Cognos gateway instance
vm-cognos1	Alias to the first Cognos dispatcher
vm-cognos2 ... vm-cognos10	Aliases to the remaining nine potential Cognos dispatchers

By default, all the aliases in the **hosts** file are mapped to **localhost**, the alias for the current machine instance.

From the Cognos Configuration tool, on the Environment tab, enter the settings listed in Table 5.2.

Table 5.2: Settings for the Environment Tab

Gateway URI	Replace "localhost" with "vm-gateway"
Dispatcher URI	Enter the ten dispatchers: http://vm-cognos1:9080/p2pd/servlet/dispatcher/ext, http://vm-cognos2:9080/p2pd/servlet/dispatcher/ext, and so on
Controller URI for gateway	Replace "localhost" with "myCognos"
External dispatcher URI	Replace "localhost" with "myCognos"
Internal dispatcher URI	Replace "localhost" with "myCognos"
Dispatcher URI for external applications	Replace "localhost" with "myCognos"
Content Manager	Enter the ten dispatchers: http://vm-cognos1:9080/p2pd/servlet/dispatcher/ext, http://vm-cognos2:9080/p2pd/servlet/dispatcher/ext, and so on

In the URIs in the table, replace the port 9080 (the default port used by the dispatcher when installed into WebSphere Application Server) with the appropriate port number in your environment. For example, if you are using Tomcat instead of WebSphere, this might be port 9300.

These configuration settings let you scale out your cloud topology from a single-instance cluster to 12 instances, simply by starting instances and making changes to the **hosts** file.

A Cluster of Size 1

Any instance can act as a single, standalone cluster of size 1. All the installed software (Cognos BI, Apache, and DB2) runs on the single instance.

A Cluster of Size 2

For a cluster of size 2, the two instances will play the following roles:

- Instance 1: DB2 role
- Instance 2: gateway, dispatcher, and content manager role

Assign instance 1 to the DB2 role by making the changes shown in Table 5.3 to instance 1's **hosts** file.

Table 5.3: Changes to the Instance 1 Hosts File

Hostname Alias	New Setting
myCognos	Instance 1 IP
vm-db2	Instance 1 IP
vm-gateway	Instance 2 IP
vm-cognos1 ... vm-cognos10	Instance 2 IP

Assign instance 2 to the gateway, dispatcher, and content manager role by making the changes shown in Table 5.4 to instance 2's **hosts** file.

Table 5.4: Changes to the Instance 2 Hosts File

Hostname Alias	New Setting
myCognos	Instance 2 IP
vm-db2	Instance 1 IP
vm-gateway	Instance 2 IP
vm-cognos1 … vm-cognos10	Instance 2 IP

Note that changes to the **hosts** file should take effect immediately. Under most operating systems (including Windows and Linux), new and updated **hosts**

entries do not require any services to be restarted. In addition, any Cognos BI instances that are already running need not be restarted for these changes to take effect.

A Cluster of Size 3

Start your third instance, and change the **hosts** files to assign the following roles:

- Instance 1: DB2 role
- Instance 2: gateway role
- Instance 3: dispatcher/content manager role

For example, the instance 3 **hosts** file will be changed as shown in Table 5.5.

Table 5.5: Changes to the Instance 3 Hosts File

Hostname Alias	New Setting
myCognos	Instance 3 IP
vm-db2	Instance 1 IP
vm-gateway	Instance 2 IP
vm-cognos1 . . . vm-cognos10	Instance 3 IP

Scaling Out to Clusters of Sizes 4 to 12

To scale out additional instances beyond size 3, repeat the steps outlined above. Add a new instance, and attach it to the topology as an additional dispatcher by assigning the remaining **vm-cognos*n*** hostnames to the new instance's IP. For example, adding the fourth instance modifies instance 4's **hosts** file as shown in Table 5.6.

Table 5.6: Changes to the Instance 4 Hosts File

Hostname Alias	New Setting
myCognos	Instance n IP
vm-db2	Instance 1 IP
vm-gateway	Instance 2 IP
vm-cognos1	Instance 3 IP
vm-cognos2 . . . vm-cognos10	Instance 4 IP

Essentially, you are dynamically scaling out your solution by adding cloud instances to the dispatcher list. The dispatcher list acts as a load balancer, so that Cognos BI will contact the dispatchers in the order they appear within Cognos configuration.

The limit of ten dispatchers (for a total of 12 instances, in this topology) was arbitrarily set for this example. It can be lowered or raised based on your requirements.

Scaling in Your Cluster

The preceding steps describe how to grow your cluster. To shrink the cluster, perform the steps in reverse. Note that when you shut down instances, in-flight transactions might be lost ("page not found") for any users currently on that server.

Creating Snapshots Using Private Images

As you develop your solution, be sure to "snapshot" your work by creating private images in the IBM Cloud. In addition to providing a backup, private images save time by recording all the software installation and configuration details for your image.

For example, by creating a private image of your operating system with any repository changes, security/firewall settings, and some common tools, you can create a starting point for other images, without having to start from scratch each time. "Base image" snapshots are a recommended best practice.

Changes to a private image create a new private image, but only the differences between the first and second image are stored. This means that when you install, for example, Cognos BI into your base image and then create a new private image, the result is a much smaller image. The image in this case needs to persist only the newly added Cognos BI software.

Note that you will not be able to delete any intermediate private images because private images are stored only as delta changes. For example, if image C is based on image B, which is based on image A, you will not be able to delete image B because image C depends on it.

Files in the Cloud

Files in the IBM Cloud are stored in two places.

Files may be stored on the file system associated with the cloud instance. This local file system is analogous to a PC's hard drive. Files stored here should be considered temporary because each instance's file system is ephemeral. If the instance is shut down or suffers some sort of failure, information persisted on the file system will be lost. In addition, files stored within the cloud instance are not accessible to other instances.

Files may also be stored on a mounted directory based on a file system from the IBM Cloud's Storage instance. This cloud file system is analogous to a network addressable storage (NAS) unit connected to multiple PCs. Files stored here can be considered permanent, thanks to the backup and redundancy guarantees of the Storage Cloud. Files stored in this mounted directory are accessible by and sharable among all your instances. Additional backup and recovery services of your storage can be purchased separately from the IBM Smart Business Development and Test Cloud.

Files that need to be stored outside your cloud images, or shared between cloud instances, can be stored in mounted directories from the IBM Storage Cloud. In a multi-image deployment of Cognos BI, consider storing the following in mounted directories, rather than in your local instances:

- File-based data sources
- Common software
- Common deployment scripts (letting you update/tweak these scripts without modifying multiple private images)
- Pre-built **hosts** files and other configuration files (to copy into your cloud instances)
- Copies of important audit or log files (so they persist even if the originating instance is shut down)
- Database backup files

6

Performance and Scalability Best Practices

Cognos 8 and Cognos 10 Business Intelligence provide a sophisticated set of products and capabilities. Sizing the architecture can be a challenging task, requiring knowledge of the system's anticipated overall workload. As a result, expect some iteration on the architecture to customize it appropriately to your requirements. This chapter provides some general guidelines and architecture recommendations to support a large-scale deployment of Cognos BI onto the IBM Cloud.

There are many factors to consider when determining the virtual hardware specifications for your cloud instances. The combination of these factors influences the amount of load the system can handle, which may vary from time to time.

In a typical data center environment, you must plan for peak loads, so IT resources are frequently underutilized. Thanks to the dynamic nature of cloud, however, peak loads are easily accommodated as needed, so the new planning standard becomes the average capacity needs. Resources can be added dynamically when required and removed at any time when the load is

reduced. The diagrams in Figure 6.1 illustrate the difference between these two environments.

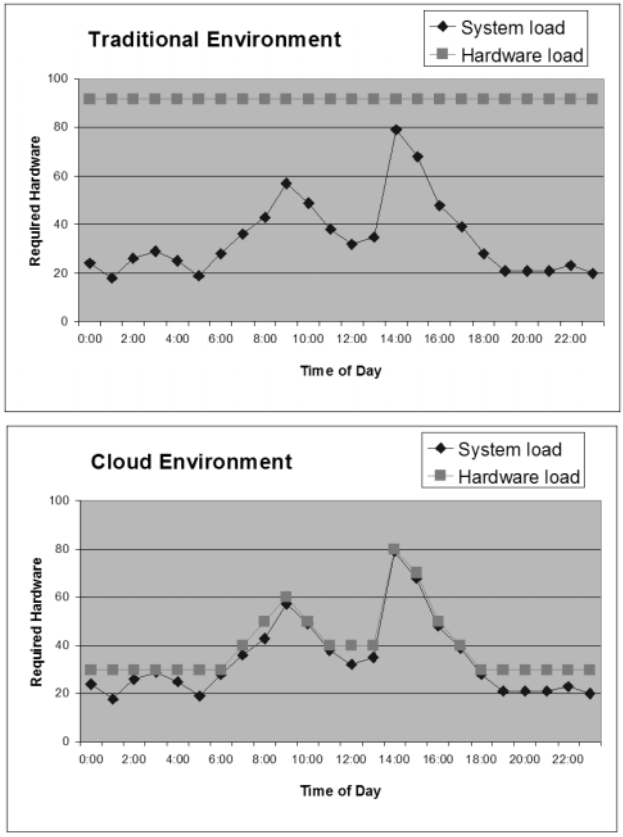

Figure 6.1: Resource requirements for a traditional versus a cloud environment.

User Community and Geographic Distribution

For the context of this discussion, the users of your Cognos solution can be categorized into three different groups:

- *Named* users are all users registered to the system.
- *Active* users are named users who are currently logged onto the system but are not necessarily in the process of submitting a request or waiting for one to finish (e.g., a user reading a report).

- *Concurrent* users are active users who are in the process of submitting a request or waiting for one to finish.

Based on IBM's experience with large-scale deployment, a commonly observed ratio among these users is 100:10:1, respectively (1 percent concurrent at any one time). That is, for each 100 named users, ten users are active, and one user is in the process of submitting or executing a request. Note that some systems require much higher or lower concurrency rates. In most cases, however, concurrency rates do not exceed 5 to 7 percent, even during peak hours. When calculating system load, you need to consider only concurrent users.

Another important factor to consider is the geographic distribution of the user community, which will also affect the average and peak load of the system.

Application Complexity

In addition to the requirements of the average number of concurrent users, the complexity of your application influences the resources required. For example, reports that involve complex database queries or advanced formatting will require more report processing capacity. This means the number of reports that can be served at any given time will be reduced. Supporting the same number of concurrent users in these types of applications requires greater system capacity.

Web-Server-Tier Performance and Scalability

For planning purposes, we suggest that an environment should be sized to accommodate 50 concurrent users per modern CPU, regardless of the users' roles. However, two factors will affect the estimation of the hardware requirement to support the average load on the Web tier. First, if SSL connectivity is used, there should be a reduction in the estimated number of supported concurrent users. More important, also consider the failover requirements of the solution. Allocate extra CPU power to fulfill your high availability requirements.

Application-Tier Performance and Scalability

Planning for report and query processing is the most important area for a highly performing Cognos solution. Report and query processing is influenced by the number of concurrent users and by the complexity of your application.

For *interactive reporting*, a general guideline is to assume a starting point of two interactive report processes per CPU, each with four report execution threads. This translates into eight concurrent, interactive reporting requests per physical CPU.

Because of the large data volumes associated with *batch report processing*, a general guideline is to assume two batch report processes per CPU, each with two report execution threads. This translates into four concurrent batch reports executed simultaneously per physical CPU.

Keep in mind that these are general assumptions based on IBM's experience, and conditions may vary from solution to solution. When calculating the required resources for a Cognos solution in the cloud, use these general assumptions, along with the expected daily average number of concurrent users, to determine the cloud hardware requirements.

Cognos BI is designed with scalability in mind. Adding more Cognos application servers during peak hours to handle heavier system loads is as straightforward as instantiating new instances of the server into the system. These newly added server instances will share the system load automatically and almost immediately. When that extra processing power is no longer required, the servers can just as easily be shut down and removed. One such method of dynamically adding and removing Cognos instances to your topology is described in Chapter 3.

Content Manager Performance and Scalability

The performance of the Content Manager depends substantially on the number of objects in a package or folder and on the security associated with the objects. For example, if a folder or package contains a large number of objects, user permissions for each object must be verified when a user with limited privileges accesses, to ensure that security rules are enforced. If there were fewer objects in the folder, or if the entire folder was not secured, far fewer resources would be required.

IBM's experience with average usage patterns suggests that one Content Manager CPU should be allocated for every four report processing CPUs. However, if your application requires more Content Manager processing capacity, then doubling the Content Manager CPU allocation would be reasonable.

Finally, because 32-bit JVMs are limited to a 2 GB memory addressing space, IBM recommends the Cognos Content Manager be deployed onto a

64-bit operating system. A 64-bit configuration is strongly advised for large-scale deployments.

Post-Deployment Consideration

These recommendations are intended to provide only a general guideline, based on our experience with average usage patterns. System monitoring is vital as resources are deployed, and it often leads to configuration adjustments. Additional resources might be required, depending on failover and load-balancing requirements.

7

High Availability Best Practices

The goal of high availability is to provide a backup mechanism that transfers requests and data processing to a standby system if the primary system fails. However, user and system requirements vary, so there is no one ideal high availability configuration. Carefully consider the settings based on your system's performance, availability, scalability, and reliability.

This chapter provides recommendations for setting up and maintaining the Cognos solution for high availability and disaster recovery, including the following:

- Cognos Gateways and Cognos Application Servers
- Active and Standby Cognos Content Manager
- IBM DB2 High Availability and Disaster Recovery (HADR)

Cognos Gateways and Application Servers

To achieve failover support at the gateway layer, install multiple Cognos Gateways onto a Web farm, one on each Web server. The Web farm entry point (commonly a router or reverse proxy server) should be able to re-route requests

to the next available Web server in the event of a Web server failure, as shown in Figure 7.1.

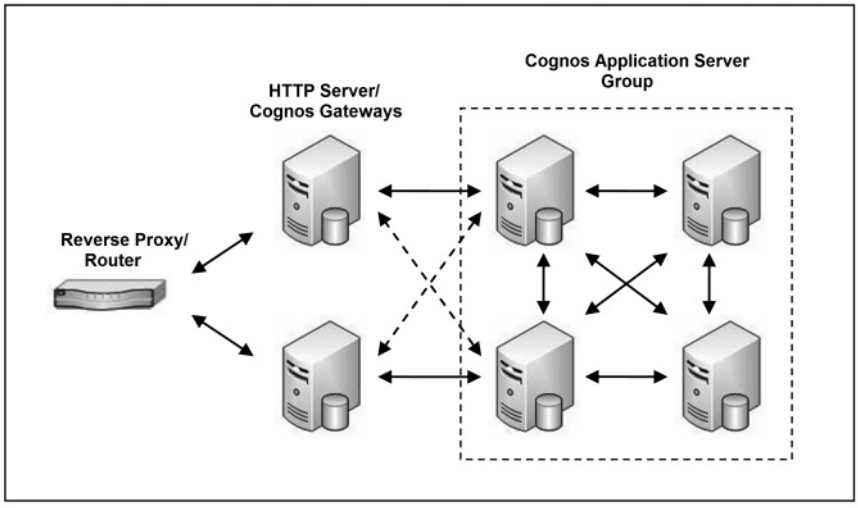

Figure 7.1: A high availability Cognos cloud environment, with a reverse proxy router.

IBM recommends that each Cognos Gateway be configured with multiple Cognos Application Servers. Requests to a gateway are routed to the first available server in the list. If that server is not available, the gateway re-routes the request to the next available server.

Note that the status of the primary Cognos Application Server for each Cognos Gateway is monitored by the gateway itself. Requests will be routed back to the primary server as soon as it returns to service.

The Cognos Application Server as a Gateway

Users who don't require gateway support for other C/C++ applications may elect to replace the Cognos Gateways with Cognos Application Servers (with all service disabled), as shown in Figure 7.2.

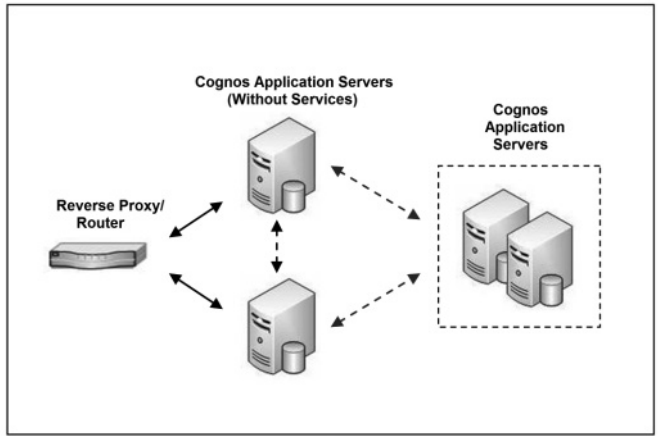

Figure 7.2: Using Cognos Application Servers as a gateway.

This topology eliminates the need to manage and maintain the configuration between gateways and application servers. That configuration is managed by the auto-service discovery feature provided by Cognos Application Server.

Active and Standby Cognos Content Manager

The failover support of Cognos Content Manager allows multiple installations of Cognos Content Manager onto the Cognos solution. Select one of the installations as the *active* Cognos Content Manager, and all other installations will be run in *standby* mode, as shown in Figure 7.3.

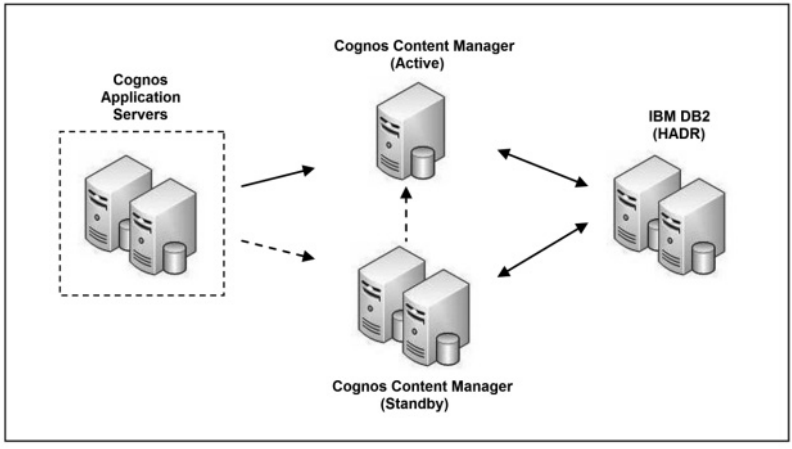

Figure 7.3: Active and standby Cognos Content Manager.

If the active Cognos Content Manager fails, the Cognos Application Server can no longer communicate with it. In this situation, the Cognos Application Server will select a standby Cognos Content Manager, which becomes the new active Content Manager. All requests will then be directed to this newly active Content Manager. All other installations of Cognos Content Manager remain in standby mode, for continuing failover support.

At least one active and one standby Cognos Content Manager should be installed for failover protection. The administrator should also be aware that when the active Content Manager fails, unsaved session data is lost. Users will be prompted to log on again after the new Content Manager becomes active.

IBM DB2 High Availability and Disaster Recovery (HADR)

The DB2-HADR is an easy-to-use feature of IBM DB2. It provides a high availability solution to handle different types of repository database failures of your Cognos solution. In the DB2-HADR environment, the administrator needs to set up two DB2 databases, one primary and one standby. Transaction logs are synchronized from the primary database to the standby database automatically, as shown in Figure 7.4.

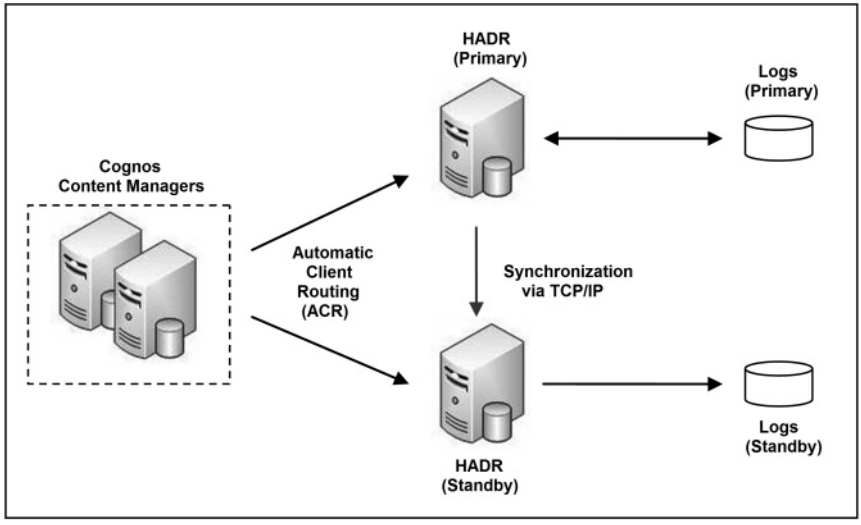

Figure 7.4: IBM DB2 High Availability and Disaster Recovery (HADR).

In the DB2-HADR environment, connections from a client to the databases are managed by the automated client routing (ACR) settings. In normal circumstances, all requests are routed to the primary database. In the event of a primary database failure, the client will receive a connection failure and will try to connect to the standby system automatically, using the information stored in the ACR settings.

The DB2-HADR environment offers several choices of synchronization modes for balancing performance, scalability, and reliability of the system:

- *Synchronous*: No possible loss of data between the primary and standby databases, but at a cost of performance on the primary.
- *Near Synchronous*: There is a remote possibility of data loss if both the primary database and standby database fail simultaneously. This mode is the best compromise between performance and reliability.
- *Asynchronous*: This is the best-performing option, but it has the potential for data loss in the event of a failure of the primary or standby instance, or of the connecting network.

The Cognos administrator should consider these choices carefully. IBM recommends that at least Near Synchronous mode be set up for all Cognos solutions and that Synchronous mode be set up for mission-critical solutions.

IBM Business Analytics and Cloud Computing

Best Practices for Deploying Cognos Business Intelligence to the IBM Cloud

This book is a practitioner's guide to deploying IBM Cognos software and the cloud. If you are a solution architect or an IT architect charged with evaluating, designing, or implementing a Cognos Business Intelligence solution to leverage a cloud infrastructure, this book is for you.

Inside, you will find pragmatic and practical information about best practices and guidelines, as well as specific software and configuration steps to help you achieve success with your Cognos BI cloud solution. Whether you are deploying Cognos 8 BI or Cognos Business Intelligence Version 10, this book will help you:

- Discover which of your analytical workloads can benefit most from cloud computing
- Learn from step-by-step instructions how to deploy IBM Cognos Business Intelligence to a cloud environment
- Leverage best practices for security, performance, scalability, and high availability

For more information, visit *www.ibm.com/cognos*.

Anant Jhingran is an IBM Fellow, VP, and CTO for IBM's Information Management division.

Stephan Jou is a Technical Architect, Research Staff Member, and Senior Technical Staff Member at IBM's Business Analytics division.

William Lee is a Senior Software Consulting Engineer at IBM.

Thanh Pham is a Solution Architect in the IBM Information Management Advanced Technology group.

Biraj Saha is an Advisory Software Developer at IBM Cognos.

MC Press Online, LLC
P.O. Box 4886
Ketchum, ID 83340-4886

PRICE: $19.95 US $21.95 CN

ISBN 978-1-58347-363-4
51995